

# It's getting better...

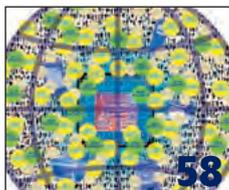
**W**hile you may not agree with the Beatles very often, in the case of Linux, I think you have to admit it's getting better all the time.

It doesn't seem that long ago that we covered the exciting new features that were due to be in the 2.4 series kernels (it was *LXF10* I believe). Now it seems that everyone is keen to see what goodies will be available in 2.6. All sorts of stuff probably, but kernel developments don't all come from the standard kernel tree. As you may have noticed, your off-the-shelf distro probably includes a patched kernel. These patches seem to be viewed with suspicion by many, but in fact, many of them are very stable, not to mention essential for certain types of kernel deployment. I would encourage you to take a look at our kernel patching feature, which not only tells you how to go about patching up your Linux powerhouse, but also details some of the most useful patches available for all sorts of users.

**Don't like your kernel? Patch it up for better security, performance and cool features**  
p44

**A fine drop of Windows API compatibility sir? WineX explained**  
p52

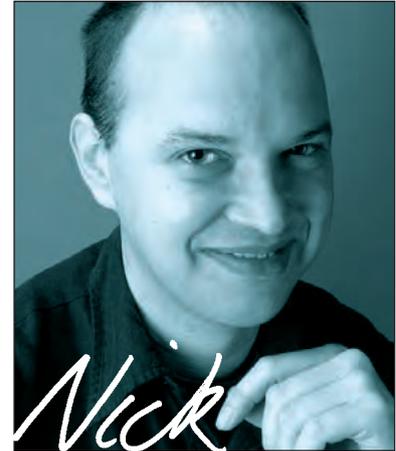
**The Internet is about to change. Find out what IPv6 will mean for you.**  
p58



*WineX* also seems to be getting better all the time, and the latest version, 2.0.1 is significantly more advanced than the initial release. Over 80 PC titles work really well with *WineX*, but the advances in Windows support have benefited applications as well as games. Support for installers, COM Objects, faster Direct2D rendering and other technologies make other applications come to life. Find out the current state of play and what the future holds for *WineX*, including an exclusive interview with founder and CEO of Transgaming Technologies, Gavriel State, in our feature on page 52.

Those of you with server interests are in for a bumper issue too, as we tackle IPv6, as well as concluding our *Apache* series with a look at scripting and PHP.

And of course, there's the roundup of astronomy software, SuSE8.0 on test, our *Linux Pro* mini magazine, A full RAD tool on the discs – it really is getting better all the time.



**Nick Veitch** EDITOR

## LINUX FORMAT

### Aims of the magazine

**Linux Format is a magazine dedicated to Linux and the Open Source community. We aim:**

- » To provide the most accurate, unbiased and up to date information on all things Linux.
- » To promote the use of Linux in business and the home, for servers and on the desktop.
- » To support the Open Source community by providing a resource of information, and a forum for debate.
- » To help all readers get more from their Linux experience by providing insightful and useful tutorials.

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



**Richard Smedley**  
We were going to take a cutting from our production editor, to increase productivity, but he objected.

**David Cartwright**  
Veteran journalist and Linux consultant, he knows his stuff when it comes to real-world Linux usage.



**David Coulson**  
Our Answers guy is a networking and security guru with plenty of sysadmin experience.

**Hoyt Duff**  
Fishing pier proprietor Hoyt spends his spare time installing Linux on anything that stays still long enough.



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

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



**Jono Bacon**  
Jono is a core KDE developer, web developer and writer. Jono is also a musician and sound engineer.

**Andrew Channelle**  
Now studying 'culture' or some such nonsense, Andy still finds plenty time to write the news!



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

**Chris Brown**  
A freelance Linux writer and Unix instructor. He has a PhD in Particle Physics, but hopes it doesn't show.

### Contact us

Letters for publication:  
[lxf.letters@futurenet.co.uk](mailto:lxf.letters@futurenet.co.uk)

Subscriptions/back issues:  
[subs@futurenet.co.uk](mailto:subs@futurenet.co.uk)

Technical help/Ask the Experts:  
[lxf.answers@futurenet.co.uk](mailto:lxf.answers@futurenet.co.uk)

Disc problems:  
[lxf.support@futurenet.co.uk](mailto:lxf.support@futurenet.co.uk)

General enquiries:  
[linuxformat@futurenet.co.uk](mailto:linuxformat@futurenet.co.uk)

Website: [www.linuxformat.co.uk](http://www.linuxformat.co.uk)

Or send your letters to:

**LINUX Format, Future Publishing, 30**

**Monmouth Street, Bath, BA1 2BW**

**Phone: 01225 442244**

**Fax: 01225 732295**

**More contact info on p113**

# Contents

# LINUX FORMAT

LXF29 July 2002

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

## Astronomy Roundup

We compare five applications which will take you to the stars. Take a break from your all-night coding session and look to the heavens.



33

## WineX 2 success

Sit back as we uncork the full story on Transgaming's successful extension of the Wine non-emulator to the DirectX8 gaming arena.



52

## What is IPV6?

There are currently not enough Internet addresses to go around. Read all about the next-generation addressing scheme which solves the problem.



58

COVER FEATURE

# PATCHING THE KERNEL

Why your distro doesn't ship with Linus's kernel, and how you can try patching too

44

Save money and subscribe to *Linux Format*. See page 96 or phone 0870 444 8645

## »» REVIEWS

**18 MORPHON XML**  
This editor provides a powerful environment for XML development, including a flexible CSS editor.



**20 CREATURES INTERNET EDITION**  
More than just cute critters. Join a thousand budding biologists with the Internet edition...

**22 SOT LINUX SERVER 2002**  
A robust server distro, with good networking and security tools, that hails from the Linux homeland.

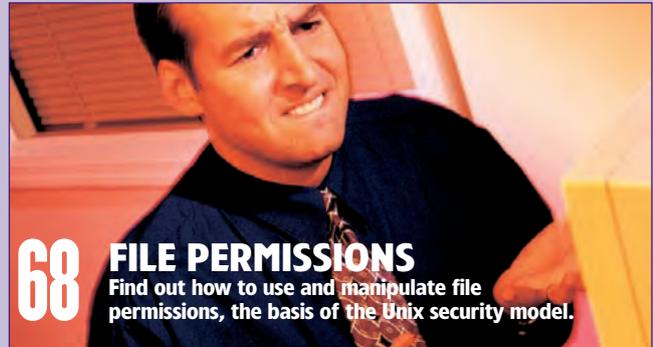
**24 SUSE 8.0**  
Find out how the latest offering from the Geeko performs as we test the new 8.0 release.

**26 ABSOFT PRO FORTRAN**  
Absoft Pro Fortran is a set of professional tools for everybody's favourite number-crunching language.

**28 LINKSCAN**  
Sophisticated link checking for your company's web site – don't be caught with embarrassing dangling links!

**30 BOOKS**  
The latest from the LXF bookshelf, including *Herdin' Cats* and *Physics for Game Developers*.

## »» TUTORIALS



**68 FILE PERMISSIONS**  
Find out how to use and manipulate file permissions, the basis of the Unix security model.

**72 APACHE**  
Follow our guide to CGI and PHP and ease web site administration by serving up dynamic content.

**76 PERL**  
Perl isn't there just for dull tasks like text processing: this issue we tackle playing MP3s.

**80 JAVA**  
So you can parse XML? What do you next? We show you have to translate your XML data into Java objects.

**82 KYLIX**  
Building on last month's installment on graphics, we get things moving with some simple animation.

**86 CGI**  
In the final part of the series on server-side programming, we investigate various methods for handling sessions.

## »» REGULARS

NEWSDESK .....	6
MAILSERVER .....	12
HOT PICKS .....	38
ANSWERS .....	90
HELP WANTED .....	108
USER GROUPS .....	110

## »» FEATURES

<b>Kernel Patching</b> .....	44
The whys and hows of kernel customising	
<b>WineX</b> .....	52
The Windows API for Linux explained.	
<b>What on Earth</b> .....	58
Is IPv6?	
<b>Emulators</b> .....	62
A round-up of 8-bit Commodores	

# Coverdiscs

A DVD or 2 CDs packed full of the latest Linux goodies **98**

**OpenOffice.org 1.0** Sun's open-source MS Office killer; **Omnis Studio 3.01** Full version of this easy-to-use RAD package; **Mozilla 1.0rc3** the monster web browser is close to release; **Netscape 7.0** Netscape's latest Mozilla-based offering; **Opera 6.0** The world's fastest browser; **Mac on Linux** Run MacOS on your PPC Linux box; **Evolution 1.0.5** Personal information manager; **Python 2.2.1** the popular scripting language.



# Newsdesk

OpenOffice.org reaches 1.0; CrossOver improvements; Schools' freedom; Palm on the Zaurus; Distros square up to competition; KDE firewall; Red Hat and Greenpeace; GNOME2 and Netscape 7 imminent.

## FLAGSHIP OFFICE SUITE

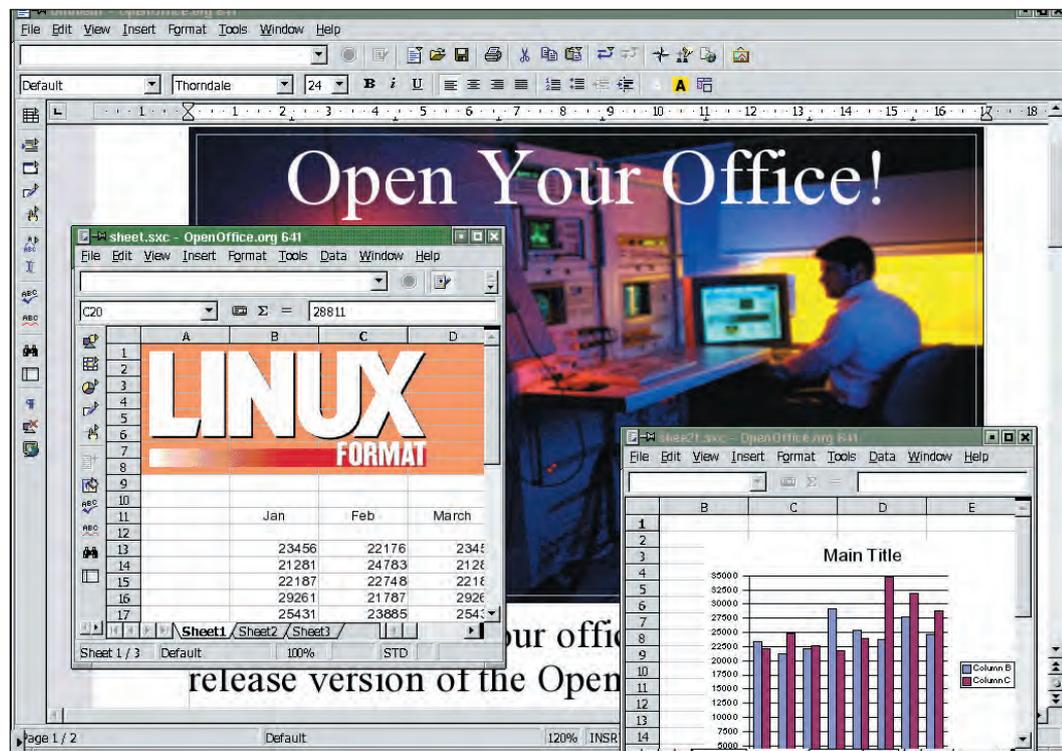
# OpenOffice.org is officially released!

OpenOffice.org (OOo), another contender for the 'biggest open source application' crown and recent winner of *LXF's* Best Business Software award, has finally hit its first 1.0 release. More importantly, it appears to have fulfilled its promise: it's stable, MS compatible, cross-platform, fully featured and free. Version 1.0 runs natively on Linux, Solaris, Windows and PPC Linux with a further 10 platforms in the pipeline.

Eric Raymond, author of *The Cathedral and The Bazaar* said the suite offered the 'single best hope' for users unable or unwilling to follow the increasingly expensive Microsoft upgrade path. "With Sun moving to a full service and support business model for *StarOffice* software, users around the globe will continue to have a free office productivity software tool through the *OpenOffice.org* open source community," he said.

Sun recently decided to begin charging for *StarOffice 6* on Linux and Windows in order to offer the sort of aftersales and support that corporate users expect, however the suite is expected to undercut the current 'default' office suite from Microsoft by 75 per cent. It remains free on Solaris.

The OOo suite consists of word processor, spreadsheet, presentation and drawing modules and includes a spellchecker and filters for MS file



**It's stable, powerful and fully functional – and even compatible with the 500lb gorilla in the market!**

formats. Project leaders say it is aimed at home users, NGOs, educators and public sector users attracted by the low total cost of ownership.

### OOo number cruncher

10,000 volunteers have contributed to the project but Kevin Hendricks, one of the key developers says there is still

a lot of work to be done. "There are many important rôles that volunteer developers can play to shape the future functionality of OOo so if you are looking for someplace to contribute, OOo can use you," he said.

4.5 million 0.x copies of the software had been downloaded when 1.0 was announced. The project is

made up of over 7.5 million lines of code – the largest open source project ever. Ximian founder Miguel de Icaza says it is also the most important, proving how capable open source software is. "People will try it, see they can get everyday work done and that open source software can work for them and be an even better solution."



Checking the small print of the EULA for that "life sentence" clause....

## MS FUD BACKFIRE

# Schooled in Linux

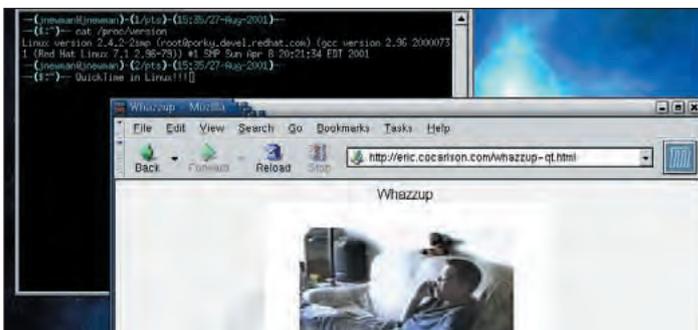
Linux companies were quick to jump on claims from Microsoft that computers donated to schools must keep the original operating system and suggesting that they should decline any offers of computers without full documentation to go with a preinstalled version of Windows. Leon Brooks of Schoolforge suggested that this policy offered an opportunity for schools and other needy organisations to upgrade to Linux. "Using Linux, OpenOffice.org and other Open Source software, a school or charity can safely accept almost any donated

computer," he said. "Simply wipe it and replace the software with Linux and Open Source applications, then use the computer as a powerful workstation or server."

After reports that some school districts in the US were facing a software audit by Microsoft (which could see them buying licenses for every computer capable of running MS software, regardless of whether it does or not) Brooks also highlighted the financial benefits of free software, which is often available at "little or no cost." [www.schoolforge.net](http://www.schoolforge.net)

## WEB COMPATIBILITY

# Crossover gets flashier

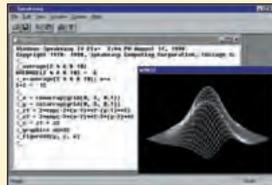


CrossOver Plugin lets you use Windows plugins in Mozilla and Konqueror.

Codeweavers have added support for Flash 6 and Superscape Viscage in the latest version of CrossOver Plugin, the software which give Linux users access to web browser plug ins developed for Windows. Also improved is integration of Freetype and Xrender,

improved MIME type associations, support for the `~/mozilla/plugins` folder in the latest Mozilla, and improvements in the uninstall procedure. There's also improved sound support and better support for multi-user installations. [www.codeweavers.com](http://www.codeweavers.com)

## NEWSBYTES



■ **Speakeasy**, which enables non-coders to 'explore the computational power of their computers,' is now available for Linux. The software's developer, Econometric Modeling & Computing Corporation, suggest the software is ideal for students, statisticians, economists and science researchers. To participate in this 'Theta' testing programme for Linux, please email [testing@emcc.com](mailto:testing@emcc.com).

■ **Netraverse** have launched version 4 of *Win4Lin*, the software that lets you run Windows on your Linux box. Enhancements include full support for WinME, support for more than 64MB RAM and access to as much virtual memory as is available in your HOME/ win directory.

■ **Gobesoft** have made a pre-alpha preview of their *GobeProductive 3.0* software available for Linux. This is the first preview of this office suite for Linux and the developers are hoping to garner impressions but are not yet ready to receive bug reports. You can preorder the suite for \$74 at [www.gobe.com/order.htm](http://www.gobe.com/order.htm) or download the evaluation version at [www.gobe.com/downloads/gobe\\_linux\\_x86\\_install.tgz](http://www.gobe.com/downloads/gobe_linux_x86_install.tgz).

■ The new **Motorola** MCF5249 Coldfire processor (designed for integrated audio and IDE controllers) is now supported by uClinux thanks to the efforts of SnapGear software wizard Greg Ungerer. "The standard ColdFire peripherals - UART, timer, and interrupts are working. I have cranked the internal PLL up to 140MHz giving something like 125 Drystone MIPS, this is a fast chip!"

■ The **Linux Expo** Birmingham has been cancelled following the demise of Sky Events. It's hoped that the idea of a Linux show in Birmingham won't die though, as a number of interested parties hope to go ahead with the event next year.

■ A letter written by Dr. Edgar David Villanueva Nuñez, Congressman of the Republica of Peru, to the General Manager of **Microsoft** gives a concise and elegant repudiation to the fear, uncertainty and doubt spread by some companies about open source software. Read it at [www.gnu.org.pe/resmsg.html](http://www.gnu.org.pe/resmsg.html) Peru is one of a number of countries evaluating Free Software in public services. In Pakistan 50,000 GNU/Linux computers are being installed in schools and universities at less than \$100 each; in Malaysia public and private enterprise task forces are investigating OSS; and in the Philippines they have the Linux School Internet Server project.

## Jono Bacon

The founder of UK Linux, KDE developer and all-round nice guy, Jono Bacon is studying at Wolverhampton University.



## COMMENT

# It's so cool that it's hot

“ Like many Linux folk, I tend to know and hang out with other Linux folk, often at user group meetings, conferences and other such events. At these occasions there is often a variety of different people, different views and different styles. Even though we are immersed in this geek culture, do you feel Linux is cool?”

As IT has continued to evolve and grow, there has been without doubt an onslaught of so called "cool" technologies; cultures that breed people with long beards, wacky t-shirts and I33t w/ys Of ty91ng. Is all of this really as cool as it is often represented; granted, "cool" is a word that is individual (just because I think sausage dogs are cool doesn't mean others must)

The one major benefit of all of these cool technologies is the excitement often expressed among cohorts of the culture. It always pleases me to meet people who truly love what they are doing, and this has been the case across many different projects and idioms.

So what would we say is the coolest of the cool? Without a doubt kernel hacking is well up there with the bearded faithful. There seems to be just something about writing a kernel patch and getting it accepted by Linus that drives so many wild. It is great to see this open and encouraging culture so rich with volunteers who love what they do.

As Linux develops we will no doubt see other technologies that are cool and motivate people to enrich and extend their skills. The great thing is that this often pays back with contributions back to the community.

Linux is cool – let's keep it that way.

## DISTRO WARS

# Gloves (and Hats) off

**T**he previously cosy world of Linux distributions has become a little harsher with the latest announcements from Caldera and Red Hat.

Caldera kicked it off with an announcement from its Caldera Global Services offshoot that they would begin offering a range of technical support packages for users of rival distros. Reg Broughton, senior vice president said the company's extensive UNIX and Linux experience and worldwide network of support centres would allow corporate users to "confidently roll out Linux solutions backed up by a worldwide professional team," regardless of the distribution they choose to use. "Caldera's history in UNIX services, our long-term Linux expertise and our global infrastructure, including multi-lingual support centers

around the world, position us better than other Linux vendor to provide business-focused services."

Red Hat, meanwhile, have begun offering a competitive upgrade rebate for users of previous versions of Red Hat 6.2+ and the latest releases from SuSE and Mandrake. Purchasers can claim back \$20 if they buy RH 7.3 Professional or \$10 if they opt for the personal edition. The offer is only open to North American, UK and German users. [www.europe.redhat.com/upgrade\\_rebate/](http://www.europe.redhat.com/upgrade_rebate/)

In breaking news though, Caldera are cooperating with SuSE, amongst others, to create a "United Linux" platform, which would go beyond the LSB to create a uniform distribution, making Linux more appealing to the enterprise sector. More info on this in the next issue!

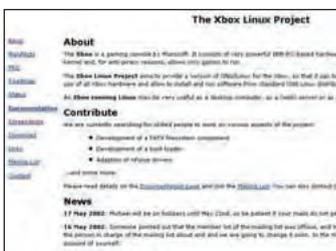


Oooh, things get competitive in the Linux world.

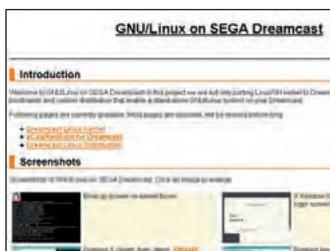
## Linux Web Watch/



PsiLinux – pocket penguin.



Xbox – obvious webserver?



Dreamcast port – it's no game!



Cash register takes Linux.

## We will make it fit!

Old computers never die they just get reborn as Linux machines.

If you have one of these devices cluttering up your loft, find it, dust it down and give it a new lease of life. Most of the projects are at various stages of incompleteness, so your experiences and feedback could help to improve their code.

With Psion withdrawing from the PDA business, now could be the perfect time for the PsiLinux ([www.psiilinux.org](http://www.psiilinux.org)) project to make an impact. The project has been running for sometime, meaning much of the work has already been done, for

instance there are tarballs available for both Revo and 5MX machines, though it takes a bit of work to get a functioning system. If you have a CompactFlash card, you can even get X11 running on a series 5.

If you've finished *Halo* and need something to occupy your Xbox, the Xbox Linux Project ([www.h07.org/xbox-linux](http://www.h07.org/xbox-linux)) could be just for you. The Xbox is just a mid-spec PC (with excellent graphics subsystem) running a stripped Win2000 kernel so it should provide plenty of hacking opportunity.

Unfortunately it's not that simple. The intention is to construct a system capable of running desktop applications, a web server or a simple and very cheap Linux cluster node.

Something that's a little more mature is the porting project on Sega's late Dreamcast console. There are a number of projects dedicated to this particular hack, but one of the best is housed at [www.m17n.org/linux-sh/dreamcast](http://www.m17n.org/linux-sh/dreamcast). The intention here is to port not just the kernel, but to create a fully-fledged bootloader and custom

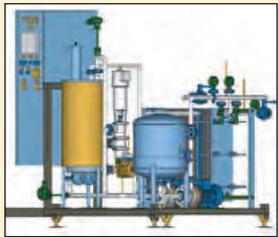
distribution. Successes so far include getting *XMAME* (arcade emulator) and *PRDoom* running on the system.

If all this is a little straight for you, Wayne Rooney managed to acquire an old cash register and thought it would be interesting to shoe horn Linux into it. He's had quite a bit of success too (<http://homepages.ihug.co.nz/~wrooney/register/>), though he did have to add memory, hard disk, floppy drive, CD-ROM, keyboard and monitor... Sort of an 'Everest' project then.

## NEWSBYTES

■ Meanwhile, **Sun** have made an enormous donation of *StarOffice 6* to the Ministry of Education in the People's Republic of China which will give millions of students unlimited access to application suite. Sun's Kim Jones said: "With this contribution of software, Sun and China's Ministry of Education will work closely with students, educators and researchers to enhance their ability to compete in a global economy, while opening the door to greater productivity and achievements throughout the Republic."

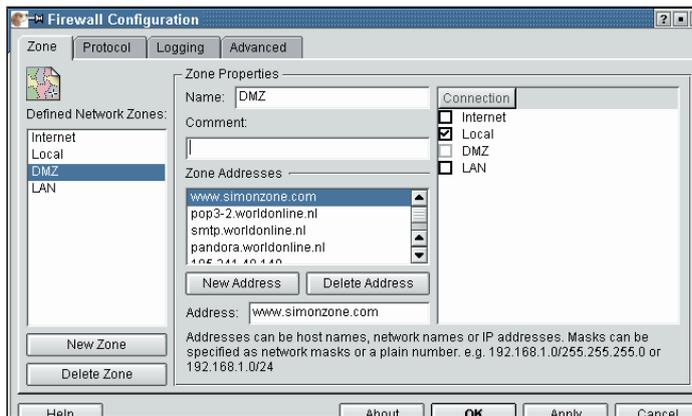
■ AOL have said that **Mozilla** will replace *Internet Explorer* in the next Mac version of their client software – this follows the announcement that the media giant's Compuserve client would also be *Gecko* based.



■ **Varicad** is a high quality, compact industrial CAD package available for both windows and Linux. Version 8.2 includes tools for 3D modeling and 2D drafting, libraries of mechanical parts (ANSI, DIN, etc), surface development, calculations of standard mechanical components and tools for working with non-graphical information. Find more at [www.varicad.com](http://www.varicad.com)

■ A recent email sent to **Shake** users has confirmed that the industrial strength film compositing tool is safe on Linux for the foreseeable future despite developer Nothing Real being purchased by Apple. Louis Cetorelli said: "In addition to OS X, Apple will continue to develop, support and sell *Shake* on Irix and Linux through the end of 2003, but the upcoming *Shake 2.5* will be the last version we release on Windows."

■ The **Association For Free Software** (AFFS) has become an associate of the Free Software Foundation Europe (FSF Europe) The AFFS is currently preparing PR material to lobby politicians, and educate the public, in the UK. Part of this effort is drafting an introduction to Free Software for non-specialists and preparing demo versions of the Debian GNU/Linux operating system for live demonstrations. Through various liaisons the AFFS is actively involved in bringing the benefits of Free Software especially to schools, other educational institutions, politicians and – of course – the public in the UK, to help create a critical mass of people using Free Software [www.affs.org.uk](http://www.affs.org.uk)



**Guarddog lets you define different security 'zones'.**

### KDE FIREWALL

## KDE gets dog standard security

### The first release candidate

(v1.9.15) of *Guarddog* firewall for KDE 2/3 is undergoing heavy testing and its developers are looking for users to test it on as many systems as possible.

Announcing the release, key hacker Simon Edwards said the software was more than just a pretty GUI masking a command line firewall. "*Guarddog* allows you to quickly and easily specify your firewall policy at a high-level, and then takes care of the rest."

In addition to testing various systems, Edwards is also keen on finding out which elements of the application are being used so he can

determine while features can be trusted. Highlights include the ability to set different 'zones' with discrete security policies, support for over 60 common network protocols, use of native *iptables* or *ipchains* packet filtering facilities, and a pretty GUI.

You can pick up RPMs for various distros at [www.simonzone.com/software/guarddog/testing.php](http://www.simonzone.com/software/guarddog/testing.php)

Also on the KDE front, version 3.0.1 of the desktop environment is out. This release was mostly about localisation, but it does include a number of bug bug fixes. You can get it from [www.kde.org](http://www.kde.org)

### PALM APPS FOR ZAURUS

## Thousands of new Zaurus apps?

### One of the early criticisms of

Sharp's Zaurus PDA has been the relative lack of applications in relation to its PalmOS/PocketPC cousins. But that problem has now been address (sort of) with the release of *QPOSE*, a version of the popular *POSE* PalmOS emulator built around the Zaurus's *QTopia* environment.

The emulator doesn't include any ROM images, which limits its use somewhat, but if you have an old Palm lying around (or sign up for the company's developer program) it shouldn't be a problem.

[www.climov.com/zaurus/qpose](http://www.climov.com/zaurus/qpose)



**QPOSE: put PalmOS on your Zaurus.**

### Hoyt Duff

The author is one of 800 Hoyts living in the USA and runs a little fishing pier when he's not dabbling with his computers.



### COMMENT

## Does it have a place?

“ The recent release of Lycoris Linux (aka Redmond Linux) follows in a line of what I often call *Brain Dead Linux* because no thinking is required to install them. From a genre of distros that include Peanut, Corel, and MaxOS, a user can simply load and go, reminiscent of a BeOS installation.

This, of course, is promoted as the future of Linux on the desktop and prompts criticism of such offerings attracting the technically challenged MS Windows user to Linux. It was this idea I wanted to write about (critical of such distros for promoting such a horrid thing) until I reflected upon it.

Linux is not a desktop OS. It is not a server OS. It's not a Unix replacement. It's not a revolution. It's not any one of those things because it's all of them and more.

Unlike any other widely used operating system, Linux is open and free in the best sense of those words and, in the Unix tradition, is a kernel plus a collection of utilities that can become greater than the sum of its parts. Because its source code is available, it can be modified, tweaked and transformed into almost anything by anyone for any purpose.

I realize the universe of Linux has room for everyone who has an interest in living in it. Just because I have no interest in a version of Linux that won't allow me to get my hands dirty doesn't mean it won't appeal to someone else. And if many computer users and businesses abandon MS Windows and use Linux, I can become wealthy writing tutorials for them as well. There is no need to be critical of others who have different ideas; we all have a place.

## RED HAT FOR MUPPETS

# RedHat goes green

**R**ed Hat launched their latest distro with a couple of high profile new customers: Kermit the frog and Greenpeace. Kermit's paymaster, Jim Henson's Creature Shop, has chosen Red Hat software to co-ordinate a number of digital production systems including its award-winning Henson Digital Performance Studio (HDPS). Henson's control systems designer Steve Rosenbluth said the upgrade – especially using Red Hat Network – has made both life and work easier. "Now we spend our time getting work done rather than trying to get our systems and applications to work," he said.

Red Hat's Mark de Visser, said the effects house was changing the way digital production was managed. "We're thrilled that the company which brought us Kermit and Ms. Piggy is using Red Hat technology."

The UK arm of Greenpeace have also switched their desktop systems to

Red Hat Linux in a bid to make their (donated) money go further. After evaluating NT, Solaris and Linux, the organisation has successfully migrated their systems to RedHat, citing not just



Red Hat: helping entertain and save the world.

cost-efficiency, but also stability and level of support. Kit Kline said that even though the organisation's IT department wasn't very familiar with Linux, the change was easy.

Finally, Reuters the global news and

information provider will soon start offering customers the option to run their market data delivery system on Linux. This move is seen as another example of the migration towards open source solutions in the finance

and banking sector. The system will be based on a combination of Red Hat software and Hewlett Packard hardware and will be worth \$200 million in sales over the next three to five years.

## BUSINESS SERVER

## Bashful server comes out of hiding

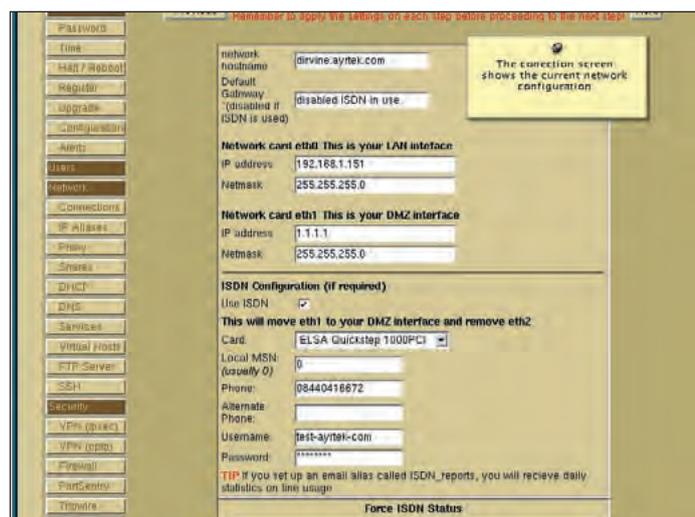
**Ayrsoft are hoping to demonstrate** the benefits of Linux to more small/medium enterprises (SMEs) with release 8.2 of *eBoxit*, a small but robust business server that has been four years in the making. Chairman D. Irvine said the company had been keeping a low profile, but that was soon to change. "We have really hidden our light under a bushel with this product, developed at great personal expense and through private investment, not many people know *eBoxit*!"

As part of this familiarity drive, user can download a free five user version of the software from Ayrsoft's website

[www.ayresoft.com](http://www.ayresoft.com)

as well as see some of the features in action through a number of online demos.

The software is designed to



Online demos provide a 'no footprint' way for evaluating the software.

allow SMEs set up and maintain a small company network including email, CRM, document management and scheduling in one easy

application suite.

As well as the free five seat version, Ayrsoft have editions available for 50 and 250 users.

## LINUX TRAINING

## Learning Linux online

### Guildford-based Merrow

Internet Services are offering a range of Linux training courses starting with the basics such as installing Linux or configuring a home network (both of which are free introductory tutorials) up to running a safe *Apache* server and programming PHP. The courses are conducted online over two 4 hour periods, and feature live contact with a tutor and the ability to remotely view his screen as he demos elements of the course.

The courses are designed to be useable over a bog-standard dial-up connection and cost £180 each. Courses in Java, C programming and Network Security are in the pipeline, and a full schedule can be found at [www.merrowinternet.com](http://www.merrowinternet.com)

## GNU DESKTOP GNOME 2

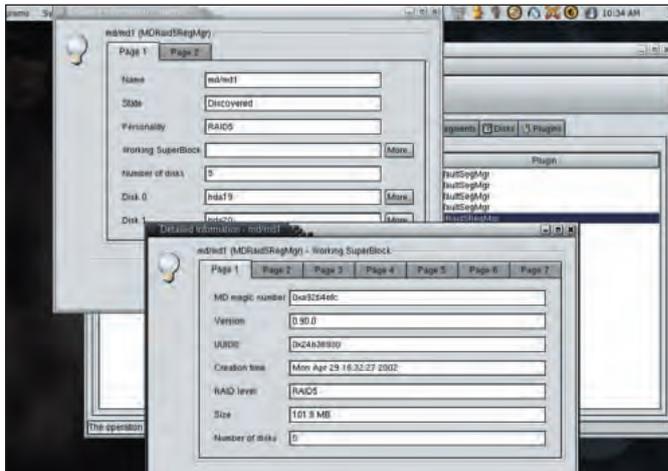
It's been three years in the making, but GNOME hackers are finally poised to unleash version 2 of the open source desktop environment on the world. Havoc Pennington said the intention was that this release wouldn't have "any new user-visible features" instead concentrating on stability and core functions. "while this was the plan and would have resulted in releasing sooner (cough), it's not really what happened."

Among more visible changes is the provision for anti-aliased fonts, new icons courtesy of *tiggert* and *jimmac* and a new default theme. One not so visible feature is the extensive use of keyboard navigation which has grown out of the GNOME Accessibility Program. The underlying system now makes it trivial for devs to implement mnemonics and tab focussing, while keybindings in general now conform to the *Swing/Motif/Windows* system.

A new help browser uses *GtkHTML2* instead of *Nautilus*.



This release is concentrating on the less visual aspect of DE design.



Manage disk, partition and volumes from one powerful tool.

### ENTERPRISE TOOLS

## World-class volume management

IBM announced a new enterprise level volume management system designed to 'streamline and enhance storage management capabilities; another key technology in making Linux the choice for enterprise level business apps.

Enterprise Volume Management System (EVMS) integrates all aspects of disk, partition, and volume management into a single design and implementation, adding industrial strength features found in proprietary volume managers to Linux. Daniel Frye, IBM's Director of Linux Technology said the system

would expand the Volume Management capabilities found in Linux today by allowing users to access data and manage volumes from virtually any OS. "EVMS technology is a quantum leap forward in readying Linux for the enterprise."

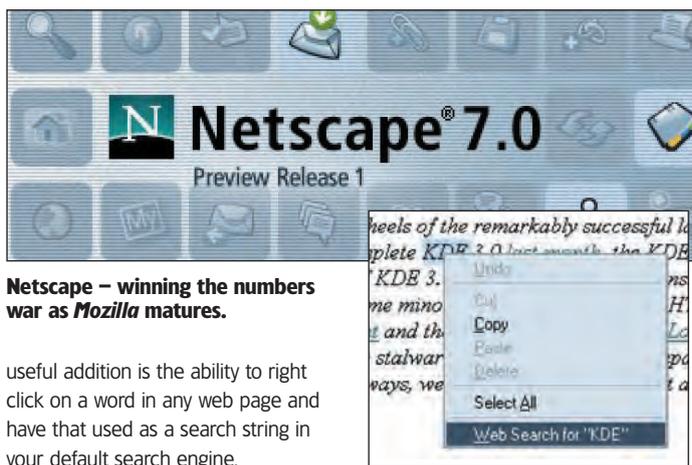
The modular nature of the app means it is capable of emulating volume managers found in other OSs. Version 1.0.0 also supports multiple disk partitioning schemes, mirroring (RAID 1), striping with and without parity (RAID 0, 4, 5), drive linking, bad block relocation, and volume groups.

[www.sf.net/projects/evms](http://www.sf.net/projects/evms)

### SEVENTH GENERATION BROWSER?

## Netscape's next gen browser suite

Netscape Communications has released a preview edition of their upcoming seventh generation application suite featuring browsing, mail, instant messaging and HTML authoring. Netscape 7 is based on the recent *Mozilla 1.0* release candidate, and is available for Linux, Windows and classic or OS-X Mac. Highlights of this release include the inclusion of tabbed browsing – which is a very popular feature in *Mozilla* – a Radio Netscape launch button, spellchecker, a download manager and integration of AOL's proprietary instant messaging service instead of *Chatzilla*. Another



Netscape – winning the numbers war as Mozilla matures.

useful addition is the ability to right click on a word in any web page and have that used as a search string in your default search engine.

### David Cartwright

David Cartwright is an IT consultant who specialises in providing Linux systems and solutions.



### COMMENT

## Where's the cache?

I've been in this industry long enough to remember when the Web was invented. In those days, Internet connections, particularly always-on leased lines, cost an arm and a leg, and bandwidth was limited. Web caches were therefore extremely popular, as they allowed you to maximise the efficiency of your Internet link by downloading something once to the cache and having it distributed to a number of users. And of course, the Unix/Linux family, generally with *NCSA HTTPd* (the forerunner of *Apache*), was the platform of choice for such devices.

It's my belief that cacheing is about to come back into the forefront of Internet connectivity. Because it doesn't take 7 hours at 3p a minute to download any more, we don't give it the value we used to. I'm as guilty as anyone for downloading the same huge file two or three times over the course of a month because I didn't stop to think: "Am I going to need this again?" and plonk it on a CDR.

The increase in convenience of the Internet will lead inevitably to a decrease in efficient use and it's a fair bet that the infrastructure will start to creak before long. It happened when 56k modems came, then with ISDN, and it'll happen now with broadband.

So caches at the ready, those of you with several machines on your broadband connection, and all you ISPs. Get that old P-266 out of the cupboard and whack on your favourite Linux distribution. Chuck a 40GB disk in for good measure, as it'll only cost a few quid. And tick the "Web Cache" option on the installer. You'll thank yourself later.



# Mailserv

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

## ★ Letter of the month

This month's winner receives a copy of **Mandrake Gaming – The Sims**

### Feature requests

I will try to keep this short so hopefully it won't turn into a nonsense rant. Some articles I would enjoy reading would be:

**Ruby** - fantastic language which makes you feel like you're sculpting your program. Already big in Japan (more popular than Python apparently), so why not the UK next?

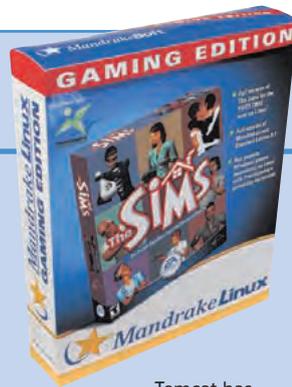
**Gentoo** - the distro, not the file manager. I realise that you ran this last year, but since it has reached a ripe old age of 1.1a maybe a revisit

is in order?

**GPG** - is PGP dead? What's going on with Mr Zimmermann? Should we be looking at OpenPGP and associated alternatives?

**Tomcat** - to follow on from your Apache series. This Java server thingy could definitely do with the LXF treatment.

**Jon Lim**, via email  
It's always nice to get some suggestions. Ruby we've been thinking about and it should appear soon. GPG was sort of covered in the encryption feature (LXF21) and



Tomcat has been covered in our Java series, though we may look at JBOSS. Thanks for your suggestions - keep them coming.

In the meantime, for your constructive comments, you can carry away our very last copy of Mandrake Gaming Edition/The Sims.

supposed to love using HTML for everything, but what works in one medium doesn't necessarily work in another.

HTML is a language for the World Wide Web and a language that has been misused. The idea of HTML is simple: to structure textual content; this is a paragraph, this is a header, and this is the body of the document.

What is being indirectly argued with the use HTML email is that there is a lack of structure to normal email, so we need HTML to provide that structure. Anyone who has needed to track an email back to its source knows there's a structure in place already. So what is HTML supposed to add to email? What benefits are we getting as a result (and the obvious cost)?

## Used Laptops

I'm pretty new to Linux, I've bought Red Hat 7 and Suse 7 only to mess with them for a very short period of time. I'm planning on going to college in about five months and want to get a laptop, used, preferably anywhere from 500-750MHz. I've not seen any ads in your magazine for Linux laptops. Is this because it's not been done yet, or do you just not have any companies selling them? Would it be possible for me to buy a laptop, install SuSE (without dual booting) and have only minimal, possibly no problems with software/hardware compatibility? And do you know of any specific brands/models of laptops that you could recommend? Thanks for your time and any help you can provide.

**Aaron Rowzee**, Texas, USA

If you are buying a used laptop, you will probably need to steer clear of two areas of possible problems - being forced to buy one with Windows installed and getting one with new

hardware that isn't supported.

Laptops are generally very well supported, especially IBM, DELL and Compaq series amongst others. I would suggest you cross reference the selection locally with the lists on [www.linux-laptop.net](http://www.linux-laptop.net) to check for compatibility. Usually the problems are with graphics cards and internal modems, but as I said, support is pretty good these days

## Cooking a new link

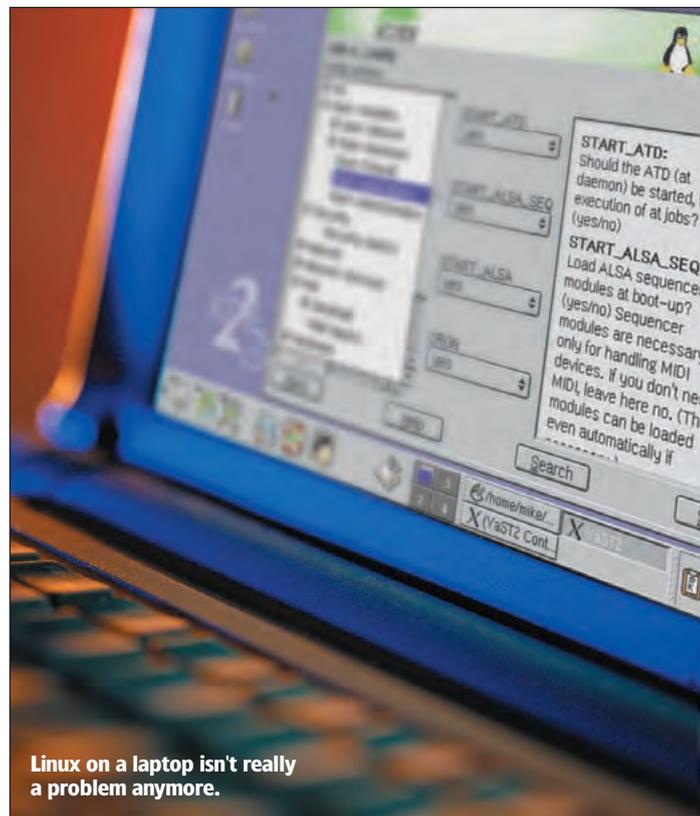
I just wanted to correct a link under the book review "The Linux Cookbook" where the link to the book's homepage is

<http://dsl.org/cookbook/>. A great magazine, keep up the good work!  
**Erik Dahl**, Sweden

Ooops. Thanks for correcting that.

## More Hoyt mail

**Hoyt Duff's** comment column about HTML email (*Linux Format* April 2002, p9) raises a good question, but falls short of being convincing. Being a web-developer I'm



Linux on a laptop isn't really a problem anymore.



So far Hoyt's column reveals bandwidth and resources as the argument against HTML emails. Yes, it takes more bandwidth (and download time), HTML email versus the plain text equivalent is a no contest (not to mention those email clients that send both versions as one email using a multi-part mime type). What have you gained? Is your message that much better because you now have words emphasised or defined as a second level-header?

The other resource that is not mentioned is the time taken for the email author to HTMLise his content. To structure textual content well takes time, it's not as natural as typing as you speak or think.

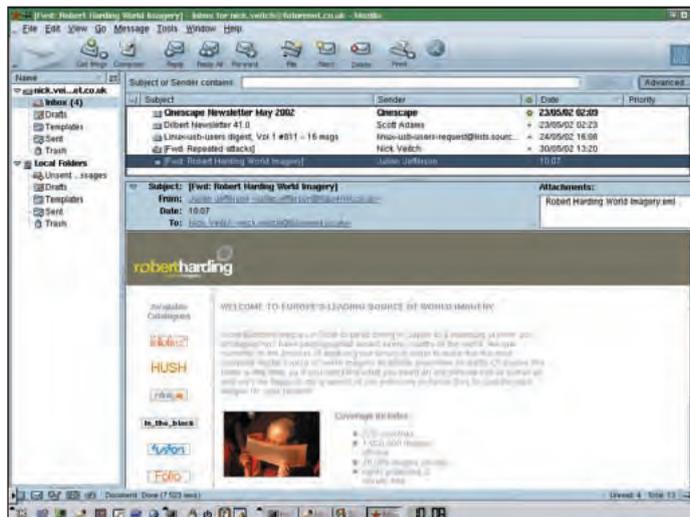
This leads to a loss in productivity since now everyone has to be at least an amateur HTML author. There is a loss in productivity on the recipient side too because of the strong tendency of HTML amateurs to create the most unreadable of pages. With HTML email you've lost the immediacy; instead of being an informal medium, it now becomes a dissertation.

If there is a need to structure the content of an email, it's also good at this point to ask why use HTML? HTML is dead because it remains an end-product, rather than a catalyst for information. It would be better to use a more natural way of structuring text, especially in a world of ever-networkable devices. An XML-based structure such as DocBook can do much more in capturing the structure people expect to achieve with HTML, with the benefit of making the document portable and in a state to transform or incorporate directly into a personal or company-wide information system. With XML we are looking toward a future in email, rather than making email look like a website.

HTML in an email is a tradeoff between download time and impact. I'd rather have a quicker download, and if someone can't get their point across in plain text, I doubt that the addition of HTML tags will make it more convincing.

Michael Davies

Croydon, Surrey



#### Shouldn't your mail-reader handle HTML messages?

All very good arguments. Personally, I can't stand HTML mail, but it is used very commonly, so shouldn't your email client support it (whether or not you want to use it)?

#### HP laptops

Thank you for the Mag, I have recently changed my subscription from another Linux magazine, and I am enjoying the difference. My reason for writing is the articles on HP and Linux. I bought a HP Pavilion laptop N5412L about two months ago and was not happy with the XP OS installed due to its carefree attitude for using space and memory. I rang HP and asked if I could get information for installing Linux, and their comments were that there was no OS other than XP that would run on the laptop. I checked Linux-for-laptops and also the HP site and not only will it not install Linux but all flavours of Windows 95-Me.

Bob White, *via email*

I imagine what the person at HP meant was that they would only SUPPORT Windows XP on the notebook, not that you couldn't run other OSes on it. It's a bit ironic as they market the Pavillion series as being tailored to the individual customer's need...

A quick look at the information on HP's site reveals that all the mentioned hardware is supported by current Linux software, although it neglects to mention the chipset, etc. I can't find any reference to anyone installing Linux on this, successfully or otherwise, but your best bet would be trying to install a modern distro like

SuSE 8.0 or Mandrake 8.2 (on our CDs and DVD last issue).

#### How dare we?

I read with interest the letter from Steve Townsley in LXF27 and his frustration at failure to install qmail. I find your answer just what I would expect to hear from Bill Gates and of the same level of arrogance. If Linux is to succeed, an indulgent attitude to novices, (as I certainly am), would be helpful. How are we, if the software is not reviewed, to know if it's good or bad? Perhaps

**"I find your answer just what I would expect to hear from Bill Gates and of the same level of arrogance. A more indulgent attitude would be helpful"**

LXF has reviewed it before I became a subscriber in which case I may owe you an apology.

Steve also makes, what I consider to be, an extremely valid point regarding the relative ease of Windows software/driver installation. I recently installed Mandrake 8.1 (many CD set) and all was okay except for me having a VANTA graphics card. So I downloaded the RPM and README from [www.vanta.com](http://www.vanta.com). I carried out the install instructions and modified the conf file (By the way, the README is 34 pages!).

After four hours of hacking I gave up and loaded Mandrake 8.1

from LXF27 DVD. Success! But Windows for all its failings would have done it in seconds and quite probably successfully.

For my sins I run a small business designing and producing electronic equipment and can't see me switching to Linux because it takes so much time to configure.

And, yes, I've installed Wine but am struggling somewhat. However, the article on LSB in LXF27, is excellent. Hopefully the ideas on standards for Linux will be taken up. Soon.

The mag? Yes. I'm keeping up my subscription.

Pete Holmes, *via email*

#### And again

I have to respond in defence of Steve Townsley email to you regarding 'A Big Question' in LXF27, page 17.

Steve emphasised how difficult it is to install software on Linux. Your reply was that in short you could not see what the problem was. The main problem, and one that has been expressed quite often in your 'Mailservers' section, is that the human interface for Linux is rubbish, in comparison to Windows 95/98/Me/NT/2000. (I won't mention XP!).

Now, I have been using Linux for 3 years, and in that time the interface has come on tremendously! But, Windows still holds the ground for Jo Bloggs on the street. Even with 3 years experience, installing software is one of those things that if you could, you would rather not.

What has to be decided with Linux is, do you sell it as a 'this thing has double overhead cam, single point injection and twin spark', or, is it sold for people who 'like the colour'? If, as I have always perceived it, Linux is for the 'techies' among us then your magazine should serve those



« people. If Linux is for techie, and Jo public, then your magazine should provide supportive feedback to those who struggle and work at trying to install software into the early hours of the morning, as Steve had done. But regardless of all of the above, keep up the good work. Your magazine, like Linux has come on along way since LXF1

Adam Cusack, *Oswestry*  
Sorry, it wasn't that we don't agree with trying new software, or didn't want to help. But Steve seemed to be complaining that it was somehow the whole of "Linux's" fault that he couldn't install qmail. I don't see what's wrong in pointing out that he was already running a perfectly good mailserver, or that there are several commercial mailserver solutions that WILL install in a few easy steps, like Gordano that we mentioned, if he was having difficulty with the one he tried. It's usually the best solution if you are looking for a 'no effort' solution.

## “Has he ever tried to install software from source under Windows? It certainly isn't as easy as just typing 'setup', and you don't get a compiler”

As it happens, not only are packages of qmail available for many distros, there is also a large list of companies who provide support for qmail listed on their site, many of whom will sell versions, with full support and an 'easy' installer at reasonable prices. There are numerous sources of documentation, including a very good book (Life with Qmail by Dave Sill, available free in PDF and HTML format. I dare say you wouldn't get such a range of support

for any Windows software, without paying handsomely for it. I empathise with the problems Steve had, because I've had them myself, and so has every Linux user, but you can't really condemn the whole Linux/Free Software community because you can't install/configure a particular package.

I don't really like these kind of comparisons because often people write in complaining about free software, when commercial support is often available. It just simply isn't comparing like for like. Yes, installing qmail from the open source code is harder than setting up MS Exchange (probably, though I've never done it). The difference is, qmail doesn't cost you \$700 (and that's only if you already have the \$1199 Windows 2000 server). Even commercial supported versions start at about \$50.

Yes, I'm sure we all wish that Linux software was easier to install, and with advances to rpm, developer tools like

Autoconf, etc, it is a lot easier to install software than it ever has been. Adoption of systems like the LSB mentioned in LXF27 will also go some way towards this, and it would be really cool if Loki's work on their installer tool wasn't forgotten. I'm not saying that this wouldn't be good. I think you'd find if you tried to install source code packages on Windows you'd have much the same problems (even more so because MS don't actually provide you with a compiler,



**Bootable DVDs or ISO images? We want to do what's best for you!**

or any equivalents of many of the GNU tools).

## On the other hand

In response to Mr Townsley's letter of LXF27, I would suggest that the many problems he faced whilst installing QMail was because he was installing from source. Has he ever tried to install software from source in Windows? It certainly isn't as simple as typing setup - and you are unlikely to get a free compiler for Windows as standard.

I would say to Mr Townsley that Linux does have many common standards, they are simply not those that he was expecting. If he wants binary-only programs, may I suggest that he sticks to rpmfind.net?

On a completely different note, I was surprised to see that you put Mandrake 8.2 on the cover DVD as DVD-bootable only, without supplying the three ISO files as well. Surely there is enough room on a DVD for both? I appreciate that more people are buying computers with DVD drives, and that more and more of these

people are looking at Linux, but in the spirit of Linux advocacy and sharing, I like to have the ISO files so I can burn install CDs for friends. Perhaps on a future cover DVD you could provide the ISO images?

Enough of the inane ramblings for today, but to finish...Love the magazine, I just wish that there were more Linux mags in the UK (I tend to read the whole of your magazine in one go, and then wait for another month for the next issue). Oh, and I'd like to give a huge plug to one of the LXF Award's Support Resources nominees: LinuxQuestions.org - definitely the best place for friendly, Linux community help.

Keep up the superb work on magazine.

Grant Sewell (aka Thymox),  
*Plymouth*

Thanks for your contribution to the debate. As for the DVD's, we included it as bootable from the DVD because we could. As most people buying the DVD edition seem to have a DVD on the machine they use for Linux, it seemed like a good idea (plus it saves burning three CDs, etc). Including the ISOs as well would have been a bit much - there wouldn't have been room for anything else (650MB x6=3900MB). Sorry.

## Xine skins

Just reading your DVD Software review and Xine "IS" skinable - even your screen captures show it!

Ok, you mightn't have FULL skinning abilities like having smaller interface, say, but that isn't what is stated. I've used Xine and it works a treat.

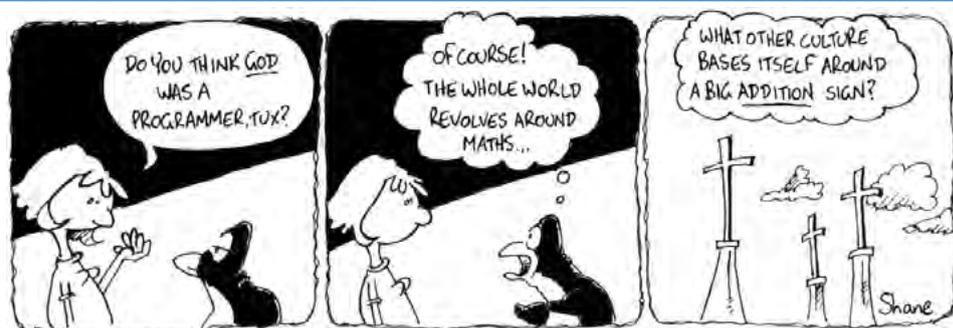
Kevin Jackson, *Manchester*

Perhaps we should define the terms more thoroughly. Anyway, thanks for pointing this out, Xine is a bit of a favourite here too.

# Helpdex

BY SHANE COLLINGE

shane\_collinge@yahoo.com





## Installing and newbies

I am constructing a dual-boot system using Windows 98SE and (once it arrives) SuSE 8.0.

Meanwhile, I am running Mandrake Linux 8.0, and I have now met the problem you have often indicated: RPM installations with missing dependencies. For veteran users, this is probably a minor hindrance, but to a newcomer from Windows like myself, it is a brick wall.

To a Windows user, a software installation binary means that the program installs and runs - although patches may later be required. This is often not the case with Linux, and my limited understanding suggests that it is particularly a problem with material disseminated by independent developers: RPM's quite often will not install and run without considerable patching for missing files, and they frequently come with either no text information on how to install and obtain the additional dependant libraries let alone information on the installation of the binary or source code itself and

what it does. These difficulties are likely to turn a complete 'Windows Novice' away from Linux - and that would be a tragedy.

RPM has been agreed as the main installation format. May I suggest the next step: universal agreement on what constitutes an installable RPM and its corresponding documentation. Source codes and tarballs are one thing - the expectation of the user is that some work on their part is going to be necessary.

However binary installation RPM's are another. I think that we need general agreement that an installable RPM should have the capability of fully installing and running without further 'messaging about', or perhaps the RPM may need to be part of an installing script plus the relevant libraries, or **MUST** carry with it a text file with full and easily understandable 'step-by-step' instructions on how to obtain the auxiliary software.

Currently, I feel that use of the *Linux Format* magazine's CD or DVD is sometimes nearly impossible for a real novice, and this is unbelievably frustrating. When you

reach the relevant directory, you find a multitude of rpms, source codes and tarballs and (from what I have seen) no text instructions on what to install, how to install, what each file does, or... whatever.

If I have missed something, my apologies. I know the magazine has notes on installations on its relevant pages, but these are easily missed and seem general purpose. I think installation notes should go with each piece of software in its directory. I would agree that a part of the problem is the multiplicity of choice that Linux offers.

May I suggest that 'long-term Linux users' recognise the above problems as real, because I am sure there are plenty of other previous Windows users (like myself) trying to make the transition to Linux who will experience these problems. Linux offers a real chance to overcome the difficulties caused by other software operating systems, but, to encourage the transition, Linux must not only be simple to install (this has now happened), but also simple to operate with respect to installations of additional software. Perhaps the journal could help considerably by having a software installation/RPM tutorial(s) covering all the likely problems.

The magazine caters very well with more advanced tutorials for Java and programming languages, but these are for people who are definitely not novices. I think the magazine could do worse than consider a monthly tutorial series for the beginner based on one of the 'starting Linux distributions' like Mandrake or SuSE because two of your stated aims are: "to promote the use of Linux in business and the home, for servers and on the desktop" and "to help all readers get more from their Linux

experience by providing insightful and useful tutorials". Novices are part of the reader group, and I know I personally would welcome the series - tutorials for the Linux novice as well as for the gurus.

And finally, I would not part with one of my *Linux Format* journals - they are forming an encyclopaedia of Linux information and are one of the most valuable and enjoyable journal sets on my bookshelves. My congratulations on a magnificent effort.

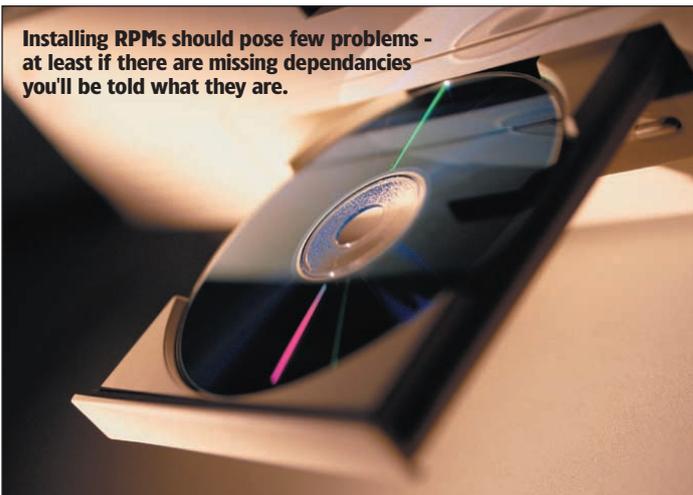
Dr Tony Young, *Australia*

Thanks very much for your input.

RPMs are usually easy to install, even if they do have dependencies. At least with an RPM, it'll tell you what files you need. With distributions like Mandrake, the software manager will even find these RPMs on the Internet and download them for you.

The problem with dependencies is more that development releases of software often require the most up to date libraries to work with - you will often find that you have suchandsuch-lib-1.03 but the software requires suchandsuch-lib-1.17 or something. Stable, release versions are not quite so demanding.

The installation instructions on the CD pages are there because for 99% of software, that's how you will install it. Any special instructions for source code is usually included by the developers. It simply isn't possible for us to create walkthrough guides for every piece of software on the disc. We are, however, drawing up a new series of tutorials which should answer newby questions about common Linux useage, including file permissions, compiling software, the kernel and so on. Hope you keep reading. [LXF](#)



Installing RPMs should pose few problems - at least if there are missing dependancies you'll be told what they are.



### 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: [lx.f.letters@futurenet.co.uk](mailto:lx.f.letters@futurenet.co.uk)

# Reviews >>>

All the latest software and hardware reviewed and rated by our experts

## LXF verdict explained

Each review is accompanied by a Linux Format Verdict to help you to assess the product at a glance (it's no substitute for actually reading the review, though). We award scores out of ten in the following categories:

**Features:** Does it provide the functions you need? Is it innovative?

**Performance:** How well does it do its job? Is it fast and reliable?

**Ease-of-use:** Is the interface well designed? Is the documentation well written, helpful?

**Value for money:** Does it have a competitive price?

For those who like numbers, the Linux Format Rating is a score out of 10 summing up the overall excellence of a product. It will usually, but need not be, an average of the above categories. We award scores as follows:



**10** The close to perfect product.



**8-9** Good, but has a few niggles.



**6-7** Does the job, but needs work.



**5-4** Average.



**1-3** An utter disaster. Back to the drawing board.

## The Top Stuff Award

If we really, really like something — we really think that a particular piece of software, hardware or any other sort of ware is the best stuff around — then we'll give it our Top Stuff Award. Only the very best will be chosen. It's not guaranteed to all products that score highly.



## THIS MONTH...

### Morphon XML

Give your productivity a boost with this powerful, visual XML editing environment. **p18**

### Creatures >>

Cook up your critters with the new and official Linux release of this online game. **p20**

### SOT Linux Server

Could this robust server distro have the tools you need? **p22**

### SuSE Linux 8.0 Professional >>

The latest edition of this popular distro put through its paces. **p24**

### Absoft Pro Fortran

Proving that Fortran isn't dead, we test a set of professional tools for Fortran development. **p26**

### Linkscan

Sophisticated link checking for your company's web site. **p28**

### Books

*Herding Cats*, *Java in a Nutshell*, *Physics for Game Developers* and *Rebel Code* **p30**



## COMING UP SOON...

### Mandrake 8.2 ProSuite

It's been downloadable for some time and we've even had it on our coverdiscs, but we've at last got our hands on the boxed version, so expect a review of this popular Linux distro next time.

### Red Hat 7.3

Can Red Hat steal back some market share from the French upstart, Mandrake?

### Win4Lin 4.0

The latest version of the software that fools your Linux box into running Windows.

### VMware ESX Server

Will the server edition of VMware's virtualization offer any real savings in manageability?

### JRun 4

Macromedia's new J2EE 1.3-compliant application server on the testbench.

## XML EDITOR

# Morphon XML Editor 2.0.4

Richard Drummond examines a visual editor for rapid producing XML documents.

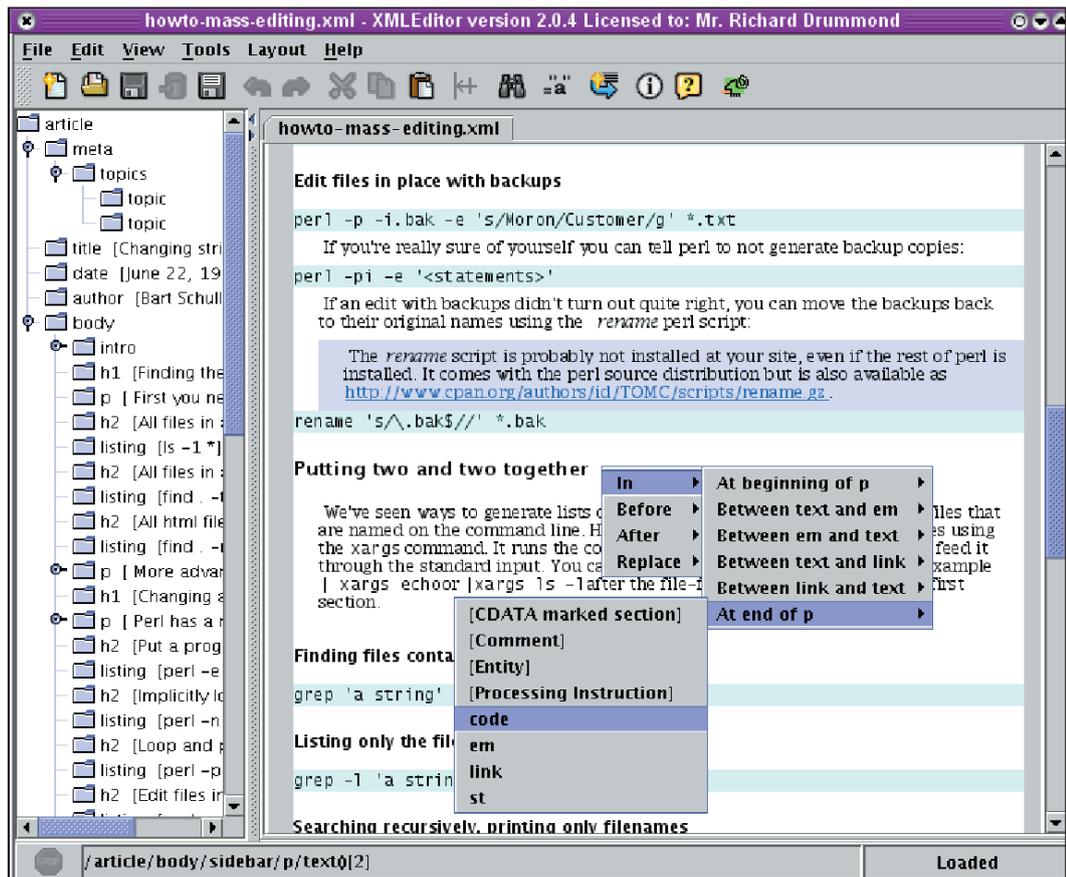
Competing products include tools bundled with IDEs like *JBuilder* and *Forté* and the Windows-only *XML Spy*.

- **DEVELOPER** Morphon Technologies
- **WEB** <http://www.morphon.com/>
- **PRICE** \$150 (Single user)

While a plain text editor is quite sufficient for knocking up the occasional XML

document, as XML becomes increasingly pervasive, the need for a more convenient solution arises. Many text editors and IDEs offer add-on modules with XML tools – such as, for instance, *Emacs's* *PSGML* mode or the Netbeans XML project – but these are often clumsy and lack transparency. Morphon Technologies' *XML Editor*, on the other hand, provides an integrated environment designed purely for the rapid production of error-free XML. The term WYSWYG has no meaning when applied to XML, but *Morphon XML* is as close as you're going to get. It makes clever use of CSS (Cascading Style Sheets) to present an XML document in the editor, thus enabling you to edit it visually.

Like many XML tools, *Morphon XML Editor* is a pure Java application – with a *Swing*-based user interface. You can download a version bundled with the requisite Java run-time, or – if you



The editor lets you pop up a context-sensitive menu with legal elements to insert, rather like code-completion.

already have a JVM – opt for a smaller download *sans* Java. The downloaded *XML Editor* requires a keyfile to be activated: you can either purchase one online or request a 30-day evaluation licence. Installation is a simple, since the package makes use of the slick

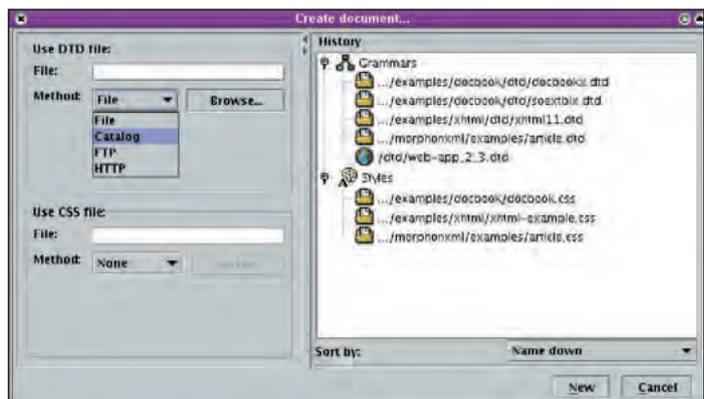
*InstallAnywhere* Java install wizard. This can automatically detect where your JVM lives on your filesystem, and perform root and non-root installations. The *Morphon XML Editor* package includes their *CSS Editor*. This integrates with the XML Editor and can be used to create stylesheets to define how your documents will look in the editor – even if you don't plan to use CSS in production. The *CSS Editor* is also available as a standalone product and is great for visually editing CSS stylesheets in general – for web applications, for example.

## Grammar class

*Morphon XML Editor* is a validating editor: it does not allow you to create invalid XML code. When you create a new document, you must specify what DTD (Document Type Descriptor) to apply and what the root element of the document is. The editor will use

the grammar specified by the DTD and then only let you add or insert syntactically correct elements. The DTD can be loaded from a local file or remotely – as can the document to be edited. When you create a new document, you may optionally specify a stylesheet for that document, too.

A neat feature of the 'Create document' wizard is that it keeps a history of DTDs and stylesheets that you have used, so that you don't have to go to the trouble of specifying their locations twice. Also handy is the fact that, if a DTD references an external DTD which the system cannot find, then *Morphon XML* will pop up a file dialog to ask you to locate it – rather than just giving up with an error. *Morphon XML* ships with DTDs and sample stylesheets for producing XHTML and Docbook documents, which will probably be the most common applications.



The history feature of the new document wizard saves time when locating DTDs and stylesheets.

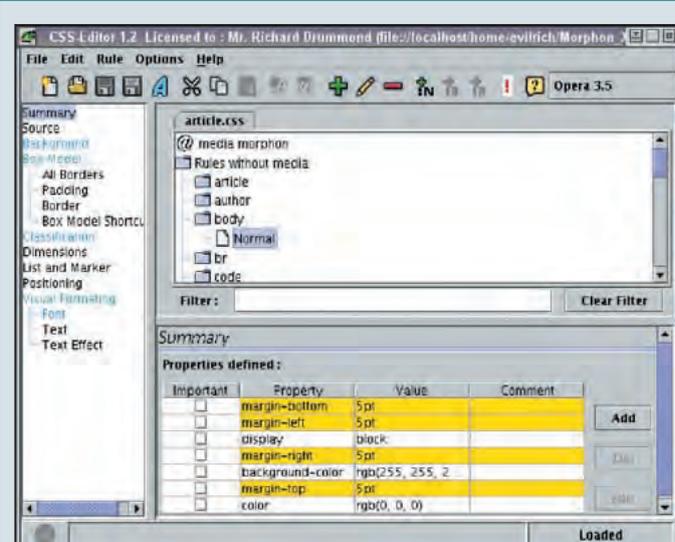
The editor is flexible with regards to interface layout, but the two main components are the tree pane and the editor pane. The tree pane provides a familiar tree based view of the elements in your document and the editor pane provides the visual view where you can edit it your document more or less as if you were using a word processor.

## Inserting elements

You can take several approaches to add or insert elements to a document, depending on what is most appropriate for the task in hand. The list of elements that are valid in a particular context can be displayed in a dialogue or as a popup menu. (The dialogue can also be permanently embedded in the main window, if you prefer, and another pane can be added to allow the setting of any of the current element's attributes and values.) The popup menu is the least intrusive method, but if you are using a complicated DTD, such as the full Docbook specification, then this becomes unwieldy because the number of valid elements may be too large to fit on screen at once. If you right click on the an element in the tree and select the 'Insert Element' option than this will either pop up the dialog or, if you have opted for the popup menu mode, it will list the valid elements in a series of sub-menus. These are also available from within the text editor pane via a right-click. When inserting an element, you must first select where in the document hierarchy the new element will be inserted relative to the current or selected element. To insert a child

## A matter of style

### CSS1, CSS2 and browser specific "standards"



The *CSS Editor* bundled with the *XML Editor* is a powerful product in its own right. It provides a point-and-click interface for building the hierarchy of rules in a stylesheet and manipulating the style properties to apply for each rule. It looks fairly daunting at first sight, but the online documentation is helpful, and the learning curve quickly levels out.

The *CSS Editor* is intended for creating the stylesheet that will determine how your documents are presented visually in the *XML Editor* itself, but it is a capable tool for general stylesheet editing. A neat

feature is that you can choose which sub-set of the stylesheet specification you want to use.

The full CSS1 and CSS2 specs and CSS Mobile are supported, but you can opt to list only those properties supported in a particular browser (profiles for the various versions of Microsoft *Internet Explorer*, *Netscape*, *Opera* and *Mozilla* are included).

If you use any properties in your stylesheet that are unsupported in the selected profile, then these will be flagged in yellow. This is a great feature for web developers beleaguered by the vagaries of browser implementation.

element, select 'in' and the position relative to any current children; to insert at the same level as the current element or it's parent, select 'before' or 'after' and, again, position relative to

the other elements at that level in the tree. When using popups, rather than the dialog, the same functionality is available with a series of context-sensitive sub-menus. The insertion functions may be operated by keyboard shortcuts while entering text in the editor pane. This is great when you don't want the act of specifying structure to detract too much from the your flow when entering text.

## XML tool box

*Morphon XML* sports many of the tools you would expect from a conventional editor or text processor. For example, an undo/redo function is included, and you can copy and paste text; however, you can also copy and paste structures, too, that is, a tree of elements and their associated data (if it is valid to do so). There's also a search and replace function, and this is particular powerful. It allows regular expression matching (*awk*, *ed*, *sed* and

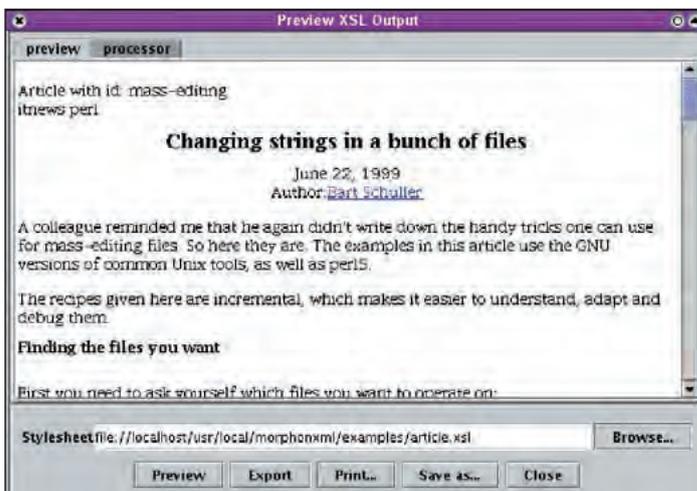
Perl type expressions are supported) and lets you search on either text, attribute name, attribute value or element name.

It also has various tools to make the job of working with XML easier. Thus there are dialogues for creating and specifying XML entities, XML notations, and a dialogue for editing the document's prologue. Tools for managing XML namespaces are included, too. If you intend to process your documents with XSLT, then *Morphon XML* supplies tools to help there too. You can select an XML transform to apply to a document and preview, and optionally export, the result.

## Conclusion

*Morphon XML* is an excellent product for rapidly generating valid XML documents. The editor pane borrows the metaphors of a conventional document processor and so is easy and transparent to use, and the on-line help is comprehensive. The pop-up menus for listing valid elements to be inserted is unobtrusive, but could be improved. Perhaps the developers should look at how IDE's such as Eclipse and NetBeans implement code-completion and adapt some of these ideas to the element insertion functions in *Morphon XML*. Feature-wise, the package is fairly complete, but I would like to see a spellchecker and the ability to optionally load and save compressed XML files.

*Morphon XML* is without much competition on the Linux. The handful of open source XML editors that are available currently fall way behind in terms of features. A direct competitor would be the Windows-only *XML Spy*, a more mature and broad-ranging product. *Morphon XML* is significantly cheaper, so the \$150 license fee is not as steep as it first looks. [LXF](#)



With the process dialog, you can apply an XSL transform and preview and export the result.

## LINUX Format VERDICT

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

A powerful set of tools that will boost productivity when working with XML.

## LINUX Format RATING

8/10

## INTERNET GAME

# Creatures Internet Edition

As he can't be trusted to water house plants, Paul Cavanagh wonders whether taking charge of a colony of furry little cute things is really a good idea.

Not so much a game as an artificial life engine, think of this as tamagotchi with bells on

■ **DEVELOPERS** Linux Game Publishing  
 ■ **PRICE** £20  
 ■ **WEB** <http://www.linuxgamepublishing.com>

Remember that whole tamagotchi thing? Some time prior to school children getting into trouble for beating each other up for their *Pokemon* trading cards, they were testing the patience of their teachers for interrupting lessons because they had to feed their electronic artificial pets. Then there was *Dogz* and *Catz*, little desktop petz that chased around your desktop if you threw a ball for them. You could download special treats for them, and upload images so the whole world could go "aaahhh, isn't he a wuwerly ickle kitty-witty?" *Creatures* is the natural progression for the artificial pet phenomenon. You're given a whole environment to manage, and you can populate it with loads of *Norns* – cute, friendly critters who learn from their experiences. While *Creatures* has been around for years, this Linux package brings together the latest full release, *Creatures 3*, along with *Docking Station* which gives your *Norns* some new toys to play with, and allows online play. If you want the full experience, you really will need *Creatures 3*, because the freely available *Docking Station* won't let you use the two new species: *Ettins* and *Grendels*. Furthermore, you won't have access to the range of environments and gadgets available in *Creatures 3*. *Docking Station* and *C3* can be linked together to make a complete package, or you can play either separately.

## A cracking start

You can begin the game with a couple of *Norn* eggs that will hatch into baby *Norns*, or start with an adult pair of creatures. For novices, the latter option



If anybody sends you creatures, this is where they'll end up. From here you can send them back, check their health, feed them or release them.

is probably best, because baby *Norns* need more care, and what's more you'll need to teach them to speak. *Norns* speak a simple language, which is largely used to express their desires, they might tell you that they're hungry, tired or bored, for example. You can teach baby *Norns* language by utilising a couple of machines, one will teach youngsters a full vocabulary in one go, and another (available only in *C3*) will teach them words one at a time. Once you've taught your *Norns* to speak, you should name them so that you can communicate with them directly by typing short sentences. You could ask them how they're feeling by typing 'express', or give them commands like 'Eric eat seed'. After you've got a few well-educated *Norns*, they will start to learn language from each other, and make suggestions to each other on how to resolve problems.

## Give them a hand

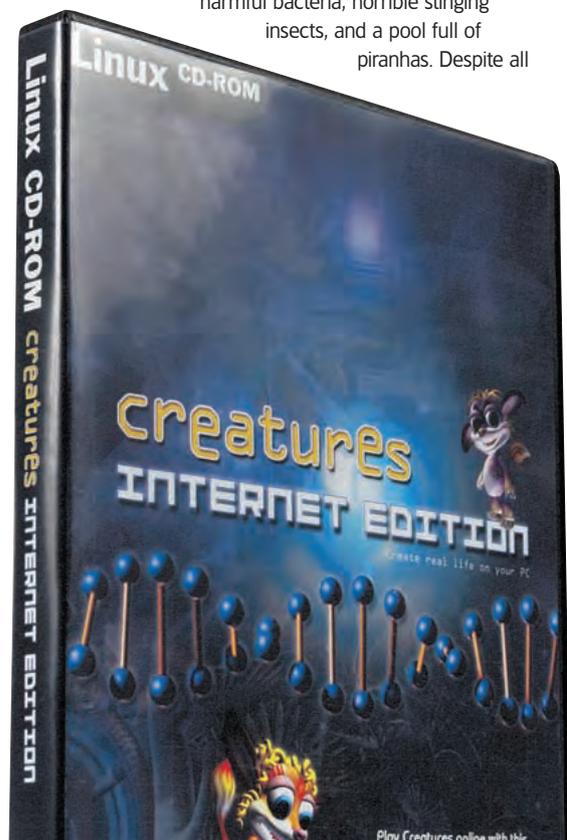
The player's presence in the game is represented by 'the hand', a cursor that is visible to the *Norns*, and can be used to lead them along, pick up objects and creatures, and administer discipline in the form of slaps, or reward them via a tickle. Words that you type in will be 'spoken' by the

hand and creatures in the vicinity will be able to hear what has been said. You can use the hand to operate machinery and gadgets too. Some of the most important machines on board the *Norn's* spaceship (drifting like Noah's ark, looking for a habitable planet) are the recycler, the creator, the replicator and the gene splicer. The creator and the replicator make tools like creature detectors, switches and sludge guns, which you can use to control the environment and protect your *Norns*. The gene splicer and medical bay are fairly complicated pieces of kit, described in the box. The recycler is essential to all of these large machines, and a few others too, because these machines run on bio-energy, a finite resource, which depletes

when you use them. In order to generate more energy you'll have to put some gadgets (or agents as they're properly called) into the recycler and convert them into energy. In *C3* standalone these machines aren't working very efficiently when you begin a game, and you'll have to find powerups, lead a *Norn* by the hand to them and encourage them to activate it. Also, you'll have to activate powerups to be able to pick up *Norns* (as opposed to holding their hands), and to be able to control the *grendels* and *ettins*. This involves leading *Norns* into some dangerous areas where they might be susceptible to disease, drowning, cold or heat. Curiously, you don't need to activate these powerups when using *Docking Station*, even if it's linked with *C3*, which is a shame because it's a good little challenge to get you into the game.

## Special species

*Creatures 3* introduces two new species, the *Grendels* and *Ettins*. *Grendels* pose a significant threat to your *Norn* population. They're aggressive, carry disease, and like to steal *Norn* eggs. *Grendels* live in a jungle area, which is swarming with harmful bacteria, horrible stinging insects, and a pool full of piranhas. Despite all



## High Tech Gizmos

Come into the lab . . .



Here are two of the sophisticated pieces of machinery to help your Norn care.

### THE MEDICAL BAY

The bridge contains loads of medical equipment, but the most advanced is the medical bay. If you lock a Norn, Ettin, or Grendel inside, you can monitor its heart

rate and energy levels. You can also analyse the levels of chemicals relating to its digestion system or state of sexual receptiveness (oestrogen, testosterone, pheromones). Toxins and bacteria are also detected here. There are hundreds of chemicals stored in this unit which



you can inject into your creatures in order to heal them, or radically alter their chemical composition.

### THE GENETIC SPLICING MACHINE

With this machine you can mix up the genes of a Grendel and an Ettin, and create a hybrid creature. Then you could

mix up this hybrid with a Norn, and you'd have a three-in-one creature. Seeing as each creature has a unique genetic code, you can create all sorts of bizarre mutations very quickly. There are two pods, put a creature in each one, hit the button, and out pops a new species.

of this, if you don't keep an eye on them, you'll often find your Norns wandering into the Grendel jungle and getting into all sorts of problems. You should keep a close eye on your Grendel population, and keep it down, and try to keep norns away from them, as Grendels are genetically programmed to behave violently as soon as they smell Norns. Ettins are less of a threat, but are prone to stealing gadgets. They were designed to be maintenance workers, and get

bored, so they'll try to occupy themselves by tinkering with technology. They live in the desert area of the ship, which along with the aquarium area, and the bridge and engineering sections completes the *Creatures 3* environments.

## GMOs

*Creatures* is much more of a simulator than a game. The only real objectives are the ones that you set for yourself. You might just want to establish a

healthy, successful Norn population – this is no easy task. You might want to selectively breed them, in order to develop Norns of a certain colour, or produce Norns of great longevity. This is entirely possible – Norns have a genetic code, containing a good deal of info about their characteristics. When a Norn is conceived its parents' genes are mixed up in a specific way, and genetic traits and mutations can be formed over time. You could just let them get on with it and see what happens a few generations down the line. The online element of the game stirs the pot up a bit – other Norn breeders can send you their creatures, which means that you can have norns that have mutated over several thousand generations turning up in your world. The creatures community is a large one, and very active – some users set up controlled experiments, carefully breeding in characteristics over a long period of time, other users are just curious about what new mutations are emerging and will allow just about any creature into their worlds. There are also plenty of agents to download into your game world. New environments, tools and toys are developed using a language called CAOS, and there's good docs and support available for people wishing to

develop for *Creatures*. Creature Labs also offer new breeds of norns for sale on their website, all of which work well with Linux. Even if you get sick of fiddling with your Norns, there are plenty of other lifeforms that inhabit the creatures world that you can play with including fish, butterflies, sharks, hawks, ants, robins, kingfishers, trees and flowers. You'll need patience and perseverance to get the most out of *Creatures*, and it can be frustrating at times. While it is a long way from being cutting edge in terms of audio and visual techniques, there is a great deal of complexity and depth here. If you're interested in genetics, breeding or biology, or if you want to develop your nurturing instincts then look no further than *Creatures Internet Edition*. [LXF](#)



You can take photographs of all of your creatures, so that you can post them on the 'net or print them out.

## LINUX Format VERDICT

Playability	6/10
Presentation	6/10
Features	7/10
Value for money	7/10

As with real pets, Norns can be frustratingly wayward little tykes, but patience and care will be rewarded.

## LINUX Format RATING

7/10

## SERVER DISTRO

# SOT Linux Server 2002

**Richard Drummond** tries out this latest professional distribution from the Linux homeland.

This distro competes against such products as Red Hat Professional and Caldera's OpenLinux Server.

- **DEVELOPER** SOT Finnish Software Engineering Ltd.
- **WEB** [www.sot.com/en/linux](http://www.sot.com/en/linux)
- **PRICE** £499.83 (including 90 days of support)

For reasons best known to themselves, Finish distro vendor SOT have reverted to the original name for their Linux distribution, previously known as Best Linux. The name SOT Linux may not be particularly pleasing to the anglophone ear, but their product – offered in Desktop and Server flavours – may well prove attractive. SOT Linux combines simple installation, a well-featured desktop, powerful configuration tools, an automatic software update system, and professional support. This month we put SOT Linux Server through its paces, and next issue we'll take a look at the Desktop edition (SOT also provide a GPL release that you can download and use for free, but no support is included).

The SOT Linux Server box comes with three CDs (install disc, extra software and source code, respectively), a chunky 900-page manual and 90-days of support. The manual is excellent. Not only does it cover installation and configuration issues, but it also features reprints of several appropriate HOWTOs to help you get *Apache*, *Samba*, *Sendmail* and so on up and running. The 90-days worth of support includes installation and configuration help via the SOT website, security announcements, automatic software updates (via the *up2date* service) and email virus protection and server monitoring services.

## Building a server

The SOT Linux Server install wizard is smooth, showing no obvious wrinkles under testing. Like most distros these



The installer and SOT Configurator have excellent tools for setting up networking and network services

days, it handles hardware detection and partition editing with non-destructive re-sizing of partitions. Only *ext2* and *Reiser* filesystems are supported, and, surprisingly for a server distro, the installer provides no tools for setting up software RAID or LVM (Logical Volume Management). However, it does have some unique and useful features. For example, when the installer probes for hardware, it gives you a list of kernel modules that it will install to drive that hardware; at this point, you can use the GUI provided to manually install any other kernel modules you desire. Interestingly, SOT Linux includes the kernel crypto modules – which can be used for such tasks as encrypting filesystems. No tools are provided to set up encrypted filesystems, however.

SOT Linux's support for networking is particularly strong. The installer makes setting up networking via Ethernet, ISDN, and ADSL – as well as conventional modem dial-up connections – simple, and support for IPv6 tunneling is included. Equally, the tools for setting up network services are just as good, and many services,

are available for those with TNT or GeForce cards.

A full-featured KDE 2.2 desktop and *Mozilla 0.9.7* is installed if you opt for the full install. The minimum install really is minimal and recommended unless you know what you're up to. You don't even get a DHCP client, which initially threw me when I discovered my network card wasn't working.

## Updater software

Keeping your system up-to-date once it is installed is the job of SOT's *up2date* software. This program checks online for updates to installed software and automatically downloads and installs them for you. Be warned, though, using *up2date* requires that you be registered for support with SOT (you won't get it as standard with the GPL release). *Up2date* lacks the sophistication of Debian's *apt*, for example, but comes with text-mode and graphical interfaces, and is adequate for the job.

For admin tasks, SOT includes an assortment of other tools, including such old friends as *linuxconf* and Red Hat's *authconfig* and *soundconfig*. SOT have provided their own GUI admin tool, the SOT Linux *Configurator*. This provides the same excellent network and network services administration tools as the installer does and also includes a function to install extra software from the CDs. It would have been nice to see a dedicated remote administration tool such as *webmin*, though. [LXF](#)

## Further information

Kernel 2.4.12  
Glibc 2.2.3  
Xfree86 4.1.0  
Apache 1.3.19  
BIND 9.1.3  
MySQL 3.23.46  
Samba 2.2.1a  
Sendmail 8.11.6

such as *Apache*, feature simple configuration dialogs to speed administration. SOT Linux features a full complement of network services, including DNS (*BIND9*), LDAP, *Apache*, *Samba* and VPN. For machines connected to both a WAN and LAN (such as routers), you may select which interfaces each service will be available on, if any.

This may be a server distro, but SOT Linux Server doesn't skimp on configuring graphical or audio hardware. The installer probes for video and sound cards, and automatically installs and configures X Windows. NVidia's proprietary drivers

## LINUX Format VERDICT

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

A well-engineered and professional Linux distro, lacking in some areas but with particularly good support for networking and security.

LINUX Format **RATING**  
**7/10**

## DISTRO

# SuSE Linux 8.0 Professional

**Mike Saunders** lifts the lid on SuSE's next gen distribution to see if a .0 release with the freshly minted KDE3 is stable enough for everyday use.

**Comprehensive distro offering latest KDE. Compare with Red Hat 7.3 and Mandrake 8.2.**

- **DEVELOPERS** SuSE Linux
- **WEB** [www.suse.co.uk](http://www.suse.co.uk)
- **PRICE** £59 (inc. VAT)

**G**erman distro vendor SuSE GmbH, sitting alongside Red Hat, Slackware and Debian on the couch of experience (having operated since 1993), has earned a healthy reputation producing value-packed boxed-sets for Linux newcomers and old-timers alike. The arrival of 8.0 is a significant event for the Linux community. So what's new? How does it compare to the other contenders? Let's take a look...

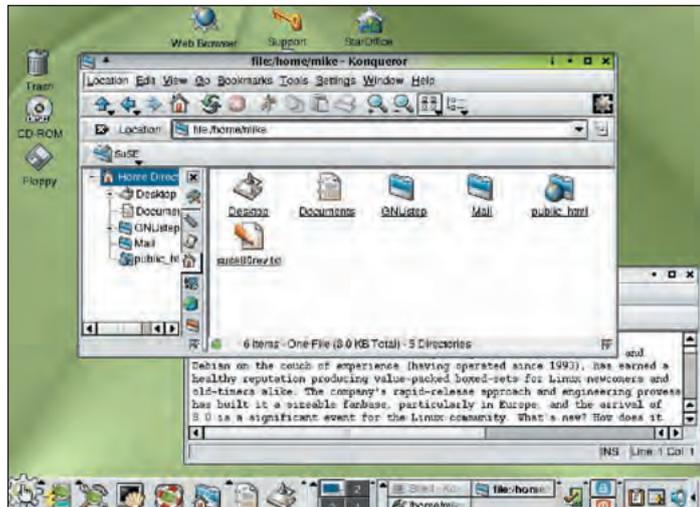
SuSE 8.0's chunky green box holds fewer surprises than its 7.x predecessors; gone are the boot floppies, fancy stickers and full-colour installation docs. Still, the three manuals (totalling 770 pages) remain the most comprehensive hardcopy guides of all the mainstream distros, with well-organised sections and copious amounts of info. Aspects covered include installation, KDE, office apps, multimedia and detailed system configuration. Typos and errors suggest a slightly hasty editing process – p339 of the *Reference Manual*, dealing with PCMCIA, claims that "information about PCMCIA can be found in the Reference Manual". Glitches aside, though, each tome is competently written and would serve as a first-rate handholder to new converts.

Seven CDs, supplied in a cardboard wallet, make up the installation, package and source set. As with previous releases, a DVD containing the entire distro is also included – a thoughtful move to eliminate disc-swapping drudgery, and an idea Red Hat has finally copied in its recent 7.3 edition.

## Installation

For this release, the archaic text-based *YaST1* has been put to rest in favour of a more flexible *YaST2*, which still supports a non-graphical console setup process. This is of particular use on older hardware, as the text-based installer will run on 48MB of RAM (as opposed to the 64MB required by its X-based equivalent). In graphical mode *YaST2* shares many similarities with its 7.x siblings, but has been overhauled and refined to produce an even simpler installation process.

Essentially, from booting the CD to sitting at a fully-fledged Linux workstation takes four steps and a minimum of interaction: choose your language, modify any automatic hardware and package selection choices you disagree with, install the software, and perform some basic post-setup tasks like adding users and configuring X. For those who require more control over



**KDE3, SuSE 8.0's default desktop, in action. For a brand new release, it's still fast and stable.**

the process, or simply prefer to do things manually, various options are available to allow for custom installs (including Braille support).

*YaST2* is a well-crafted tool which never feels too intrusive, and boasts a healthy range of features including Win9x FAT32 partition resizing, networked installs and journaling filesystems (*ReiserFS* being the default). *LILO* is written to the disk as the bootloader

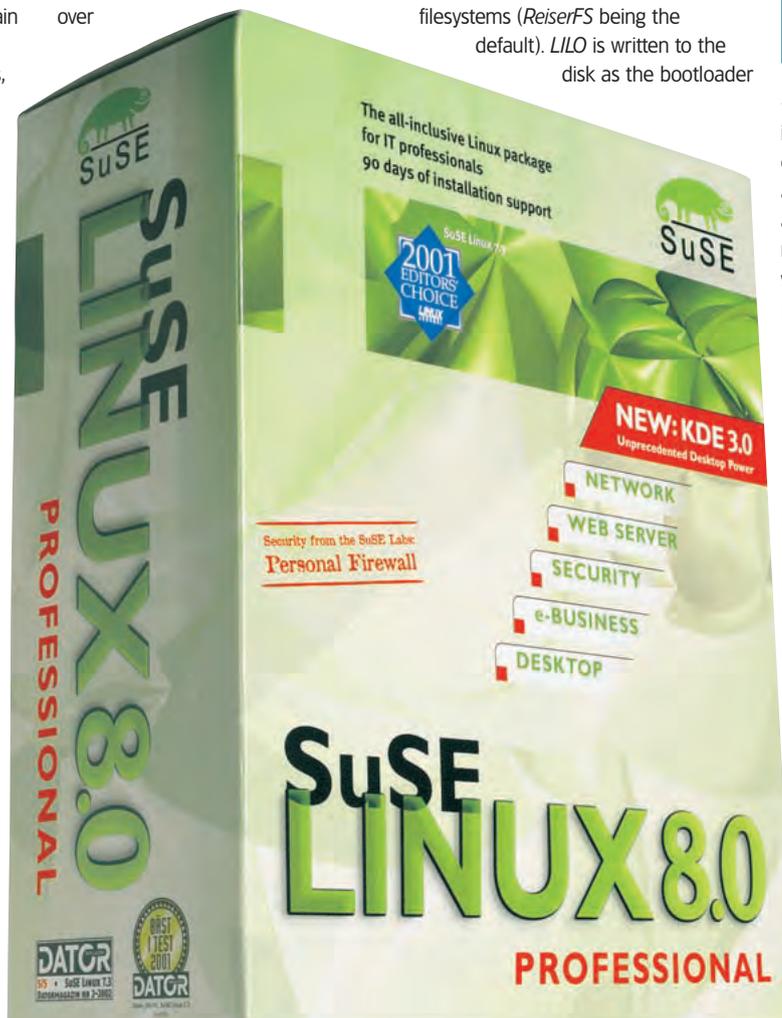
## Further info

**Installer:** Text/graphical  
**Package format:** RPM (3.0.6)  
**Kernel version:** 2.4.18  
**Glibc version:** 2.2.5  
**Default compiler:** GCC 2.95.3  
**XFree86:** 4.2.0  
**Default desktop:** KDE 3.0.0  
**Support:** 90 days installation

– we would prefer to see *GRUB*, but it's not a major concern. Hardware detection is usually satisfyingly accurate, with the mouse, sound card and X configured correctly on our two main test machines. The X setup wasn't perfectly optimal, but some post-installation tweaking corrected that, and some users have reported troubles with PCMCIA and SCSI when upgrading. In all, though, *YaST2* continues to improve as a friendly tool, while still providing scope for advanced custom setups if needed.

## SuSE 8.0 in use

With 7 CDs of software, it's no surprise that SuSE offers almost every open source app of significance in its set, along with a few commercial programs too (over 2,500 packages in total). The base system is built around kernel 2.4.18 and *glibc* 2.2.5; *XFree86* is version 4.2.0, supplemented by *Mesa* 4.0.1.



## What's new in 8.0?

### Across the board improvements

■ **KDE 3.0** – The most user-visible software increment in 8.0, KDE 3 brings its own roster of new features and goodies to the party. Built on version 3 of the Qt toolkit, this latest release of the immensely popular desktop sees improvements to *Konqueror*, supporting utilities and base libraries. SuSE has integrated its programs and docs well with KDE – the desktop has links to its support pages, the menus are suitably crammed with apps, and the scheme and theme of deep lime green is consistently present throughout. Despite the natural concerns with a .0 release, we found KDE 3 to be a stable performer in over a week of regular use.

■ **LSB compliance** – The Linux Standards Base aims to bring greater consistency to the ever-growing number of distros by establishing an accepted set of tools,

libraries and filesystem layout features that each should have. This should be a remedy to the numerous headaches commercial vendors face when porting to Linux, and should enable one package to be installed on multiple distros with less hassle. SuSE's support of LSB is an important step for the Linux world, and other vendors are expected to follow.

■ **Revamped init-scripts** – Previous releases used a large BSD-style `/etc/rc.config` file to configure the system on bootup; this has been replaced by a number of smaller text files in `/etc/sysconfig`, bringing it more in line with Red Hat. Whichever is better is a matter of personal preference for the experienced Linux user, but *YaST* makes the change mostly seamless.

■ **SuSE Firewall2** – A super-simple point-and-click tool to set up packet filtering

and secure the machine. Choose the network interface, select which daemons and services the outside world should be allowed to access, finalise any other details such as logging, and the script is written. While the truly paranoid will prefer to do things manually, it's still a highly useful util for quick, basic security.

■ **YaST2 updates and new modules** – As mentioned, the old text-based *YaST1* tool has now disappeared in favour of an ncurses front-end to *YaST2*. While text mode may seem redundant on a SuSE desktop box, for a server setup it's supremely useful – no X required, so it's better on older hardware, and it can be used to configure a system via *SSH* or *Telnet*. Also, new modules such as a Backup system have been added to *YaST2*, while the network and email modules have been improved.

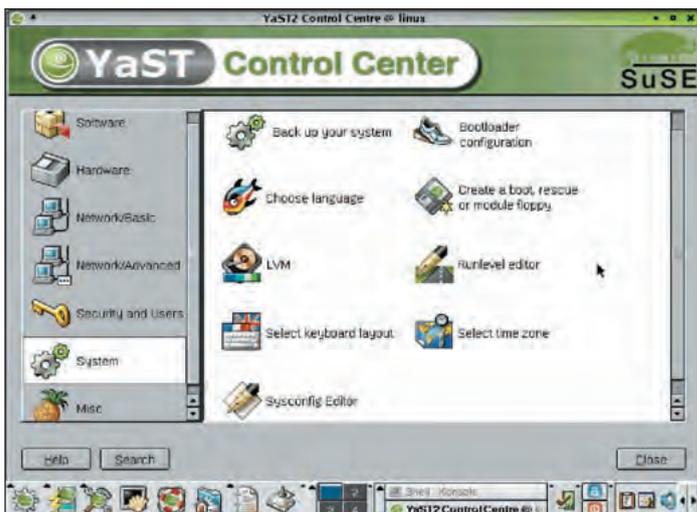
While we came across a couple of glitches in testing, it's still a mightily impressive accomplishment and remains SuSE's proudest feature.

## In conclusion

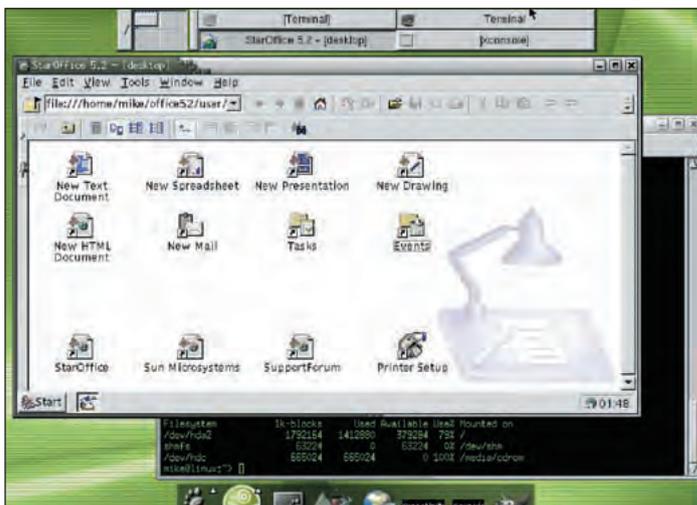
Once again, SuSE have produced an exhaustive collection of all things Linux, and at a respectable price – better value than the equivalent RH set. The 7 discs of software negate the need to spend hours hunting online, the base system is up to date, and the *YaST Online Update (YOU)* system facilitates quick and easy access to security updates and bugfixed packages. SuSE has gained a large community around the 'net, so any installation or usage niggles can be settled on websites or newsgroups if the 90 day support isn't enough.

Despite the added value of hefty handbooks and a wide range of software, those with a solid Linux background or who simply want to run a basic server won't necessarily need all the docs and packages provided; in this case, the cheaper and smaller Personal edition may be more suitable. The beefier Pro box on test here would be a better choice for a complete newcomer who wants to learn about the OS in-depth and explore a wide range of software, or the intrigued sysadmin considering deploying Linux but requiring superb documentation.

Considering the troubles other .0 releases have suffered, and that this version is particularly vulnerable with new *YaST* features and KDE 3.0 as the default desktop, SuSE are to be applauded for constructing a generally well-behaved and flexible distro. The few issues that need to be ironed out just prevent it scoring a 9, but none are critical – SuSE 8.0 Pro is still a robust, friendly, well-engineered and complete Linux distribution which makes for a particularly fine desktop OS. Recommended. [LXF](#)



**YaST2**, the central configuration point of SuSE, displaying some of its modules yesterday.



GNOME addicts need not fear, although the standard theme is quirky.

These are the latest releases at the time of writing, making for a solid and up-to-date working environment. SuSE chose not to follow Red Hat and Mandrake's lead of shipping with the GCC "2.96" snapshot as standard, instead sticking with the tried-and-tested 2.95.3 version. It uses the popular RPM package format, although installing apps directly from the discs can prove awkward because of the directory layout – *YaST* is better.

Other worthy inclusions are GNOME 1.4.1 (with a bizarre MacOSX-esque theme), *Window Maker 0.80.0*, *IceWM 1.0.9*, *Emacs 21.1*, *GIMP 1.2.3*, *StarOffice 5.2*, *Netscape 6.2.1*, *Mozilla 0.9.8*, *Apache 1.3.23*, *BIND 8.2.4/9.1.3*, *Sendmail 8.12.2*, *MySQL 3.23.48*, *PostgreSQL 7.2*, *Java2 SDK 1.3.1*, *Perl 5.6.1* and *Python 2.2* – software for every occasion, be it a multimedia desktop, development workstation or server. Puzzlingly, the *GIMP*, *Emacs* and some development packages are missing from the default install – the development libs could trouble a new user wishing to compile software from scratch.

General administration chores are best performed through *YaST2*, which remains a consistent, approachable front-end to the workings of the underlying system. With the addition of new modules and tweaks for existing components, *YaST2* is able to tackle almost all straightforward tasks, be it network setup, new hardware installation or user management.

## LINUX Format VERDICT

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

Another thorough and versatile offering from SuSE, with plenty of Linux for your money.

## LINUX Format RATING

8/10

## FORTRAN COMPILER

# Absoft Pro Fortran 7.5

Is there any scope for a commercial Fortran compiler on Linux? **Biagio Lucini** investigates Absoft's contender.

**A speedy, commercial alternative to g77.**

- **PUBLISHER** Absoft
- **TEL** +1 248 853 0050
- **UK DISTRIBUTION** For Linux
- **WEB** [www.forlinux.co.uk](http://www.forlinux.co.uk)
- **PRICE** See *Pricing* box

**N**umerical computations are an important sector of the modern computing world, driving big hardware (think of supercomputers) and software. The latter usually concerned with the creation of fast and efficient code, possibly speeding up the development by heavily recycling existing code. Among the programming languages, Fortran, with perhaps has the biggest collection of numerical routines available to programmers, is – despite being quite old – still very popular for numerical simulations.

Unfortunately native Fortran support in Linux is insufficient for many real-world applications: the *g77* compiler does not have the maturity of its cousin *gcc* and only supports the old Fortran 77 (F77) standard. For this reason, at least until *g95* (<http://g95.sourceforge.net>) is ready, the new Fortran 90 (F90) and Fortran 95 (F95) will be out of reach for the Free Software environment. Hence Fortran compilers for Linux, such as *Pro Fortran* from one of the longest running players in this arena, Absoft.

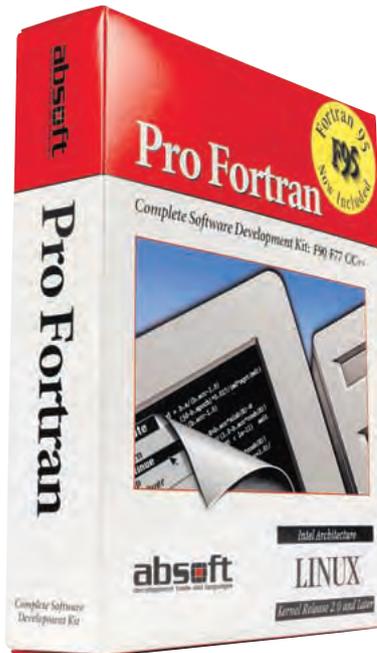
## Features

The Linux version of this product is available for the i386 and the PPC

architecture. The core of the latest release of *Pro Fortran*, 7.5, is a fully-featured compiler for F77 and F90/95. To be more precise, *Pro Fortran* comes with different executables, one for each revision of the Fortran standards to date (though the F77 compiler is not strictly needed, F90/95 being backward-compatible with F77). The F77 compiler is standard-conforming, and in addition supports many popular vendor-specific extensions of the language. This is a good feature, since usually those extensions are responsible for failure in porting code to Linux. Similar features also characterise the F90/95 compiler, the result of a five-year collaboration with the famous supercomputer vendor Cray.

The integration with the host OS is very good: *Pro Fortran* is link-compatible with *gcc*, *g77* and *f2c*, and offers familiar options to Linux developers. Highlights among the new features in the 7.5 release include new optimisers for AMD Athlon and Intel P IV processors, support for data files larger than 2GB and compliance with the Linux Standard Base.

More than just a collection of Fortran compilers, Absoft *Pro Fortran* is a good software development solution, since it can compile C and C++ code (which makes linking C/C++ subroutines a breeze), comes with many precompiled and optimised mathematical libraries like *LAPACK* (new in this release is the inclusion of the *ATLAS* libraries), the *DISLIN* graphical package and *Fx*, a graphical



debugger that supports C, F90/95, F77 and assembly. Optional software like the *IMSL* Libraries is available at additional cost. Also at extra cost, bundles for multiprocessor computers (both Beowulf clusters and SMP systems) are also available.

## Installation

The product comes in a box that contains the installation CD, a printed book containing the *Fortran User Guide* and the *Fx debugger User Guide* (the pdf version of these and other

docs are installed with the compilers) and a leaflet containing instructions for the installation (available as a README file in the CD). Of course support is available after registration.

We installed *Pro Fortran 7.5* for Linux/i386 on a P IV 1.8 GHz/256MB RAM system (see the *System Requirements* box for the minimal specifications of the hosting system). Mandrake 8.2 was chosen as one the most recent distributions at the time of writing, which allowed us to test the installation and the performance of *Pro Fortran* on an up-to-date Linux system.

The core compiler files are contained in one package provided in RPM and tar.gz format.

According to the documentation, the rpm is not relocatable, so if you want to install *Pro Fortran* on a directory other than */opt* you are forced to use the *gzipped* tar file. The rpm installed very smoothly and after setting the required environment variables (according to the excellent instructions given in the README) the compiler produced its first executable. The whole process did not take more than two minutes. We also tested the installation via the tar.gz file, and this gave some more problems, since the X default file needed by the graphical version of the debugger was not present in that package. We had better luck by converting the rpm to a tgz

```

F90/F95(1)                                     F90/F95(1)
NAME
  f90/f95 - Absoft Fortran 90/95 compiler version 3.0
SYNOPSIS
  f90/f95 [ options ] file ...
DESCRIPTION
  f90/f95 is the Absoft Fortran 90/95 compiler. It compiles
  Fortran 90/95 source code into executable programs or
  relocatable object modules.

  It compiles each file based on its extension. All files
  that end with .f, .for, .F, and .FOR are compiled as FOR-
  TRAN Fixed-Form source code. Files ending with .f and
  .FOR are first sent through the C language preprocessor,
  cpp(1), before being compiled. All files that end with
  .f90, .F90, .f95 and .F95 are processed as Fortran 90/95
  free source form. Files ending in .f90 and files ending
  with .c are compiled by the C compiler, cc(1), and .s
  files are assembled by as(1). All other files are passed
  to the linker, ld(1).
  
```

*Pro Fortran* perfectly integrates into a Linux system. Here is the *f90* manual page – available, as you would expect, for whatever Linux tool.

## Pricing

### Bundle breakdown

Product	Number of Licence	Price
Pro Fortran	1	\$899.00
Pro Fortran+IMSL	1	\$1199.00
Pro Fortran Cluster	2	\$1899.00
Pro Fortran Cluster	5	\$2695.00
Pro Fortran Cluster	10	\$4319.00

Prices for *Pro Fortran* bundles. The cluster product refers to a cluster with up to 16 nodes. Other options and academic discount are also available.

package with *alien* and then unpacking the *tgz*. For this reason, we recommend the latter procedure if the software must be installed in a directory different from */opt* or if the distribution used is not RPM-based.

One thing we did not really like was the licence management – the licence file is created by a not-so-smart executable which understands neither the backspace nor the delete key. This means every time a character is mistyped, the product code must be re-entered. All this is to create a simple text file containing just the codes entered. In our opinion this product deserves better licence management software.

## The compilers

The Absoft *Pro Fortran* compilers are command line tools. The way to invoke them will not surprise those who have used similar command line compilers (such as *gcc*). The name of the executables are *f77*, *f90* and *f95*, and they are accessible after setting the **ABSOFT** environment variable to the absoft subdirectory of the installation directory and pre-pending **\$ABSOFT/bin** to the **PATH**. The names of the executables are quite generic and may not suit people that want to use several compilers. Unfortunately the *Pro Fortran* compilers can't be renamed with the *mv* command, since for the compilers to work the names of the compiler executables must be the default one, although a makefile allows the creation of executables with

**Table 1**  
Speed statistics

	<b>f90</b>	<b>g77-2.96</b>	<b>g77-3.0</b>
<b>LQCDMC</b>	<b>16.97</b>	<b>17.92</b>	<b>17.72</b>
<b>AC</b>	<b>6.40</b>	<b>19.23</b>	<b>10.16</b>
<b>TFFT</b>	<b>32.40</b>	<b>26.61</b>	<b>25.75</b>
<b>LP</b>	<b>5.19</b>	<b>11.45</b>	<b>10.42</b>
<b>LM</b>	<b>8.88</b>	<b>8.46</b>	<b>8.30</b>

Benchmark results. The numbers give the time of execution in seconds.

## System Requirements

486DX 66 or better  
32MB RAM  
50MB HD space  
Kernel 2.0 or later  
*Glibc 2.1* or later

```

File:sun_finitet3.f Unit:main Line:10 NO PROCESS
1 C*****
2 C
3 C As sunTb12.input, but corrected to have complex line vac;
4 C also includes a simple scalar glueball ; and 1,1,1
5 C for sub
6 C
7 C If NCOL=/=3,4,5,6 then need to choose SU2 subgroups by hand
8 C in RENORMB and UPDATG
9 C*****
10 IMPLICIT REAL*8 (A-H,O-Z)
11 PARAMETER(LX1=5,LX2=12,LX3=12,LX4=12)
12 PARAMETER(LSIZEB=LX1*LX2*LX3,LSIZE=LSIZEB*LX4)
13 PARAMETER(NCOL=4,NCOL2=NCOL*NCOL)
14 PARAMETER(LMAX=LX4/2+1,MAXDLS=LMAX-1)
15 C
16 PARAMETER(ICALLG=5,NITER=50,NMERSL=NITER/ICALLG)
17 PARAMETER(IBLOK=4,IBING=5,NUMBIN=2)
18 PARAMETER(PARBS=0.25,PARBDS=0.05)
19 PARAMETER(ICALLQ=50,NCOL=20)
20 PARAMETER(QCUT=0.02,RCUT=4.0,MAX=999,MAXP=99)
21 C
22 COMMON/ARRYS/U11(NCOL2*LSIZE*4)
23 DIMENSION V11(NCOL2*LSIZE*4)
24 COMPLEX U11,V11
  
```

### The Fx main window.

the wanted names. Another drawback of the compiler is that many options have a meaningless name – e.g. to disable stack alignment you have to pass the option **-B112** instead of something like **-disable-stack-align**.

The two parameters for judging a compiler (at least as far as numerical computations are concerned) are syntactical accuracy in the compilation process and speed of the generated executable. Unfortunately the Absoft *Pro Fortran* compilers do not excel for syntactical accuracy. A test performed by Polyhedron ([www.polyhedron.co.uk](http://www.polyhedron.co.uk)) shows that the *F77* and the *F90/95* compilers have a success score respectively of 18% and 31%. What *Pro Fortran* is good at is supporting extensions to standard Fortran. This is achieved thanks to a wide range of carefully documented options.

Testing for speed of execution of the compiled code is a dungeon full of pitfalls. Again, Polyhedron maintains a table of benchmarks for different Fortran compilers. The conclusion is that *Pro Fortran* is one of the best performing, and is much better than *g77*. This conclusion will not surprise those with direct experience of past releases of *g77*. However, we decided to redo some of the benchmarks on our test system (the “official” ones having been executed on a PIII) and to add another one to the bunch. Those tests are a subset of frequently (ab)used numerical technique: a Lattice QCD Monte Carlo (LQCDMC),

employed in Theoretical Particle Physics; the computation of correlation of pseudo-random numbers (RAN); fast Fourier transforms (TFFT), used in signal processing; the dual precision linpack benchmark (LP); and large matrix multiplication (LM), which is at the basis of many numerical problems. Compilation commands we used are: **f90 -s -O2 -B100 -xINTEGER file.f** respectively for *Absoft f90* and *g77*.

Both *gcc-2.96* and *gcc-3.0.4* have been tested, and will be referenced respectively as *g77-2.96* and *g77-3.0*. We wanted to test *Absoft f77* as well, but it failed in compiling our codes. This is not a big problem, since *F90/95* are backward-compatible with *F77*, but shows somehow the limits of the *f77* product.

The results, described in **Table 1**, should be taken just as an indication. With this caveat, what they seem to show is that generally speaking Absoft has shipped a good product that performs pretty well. However, it does not always take a clear lead over *g77*, so whether your investment is worth it depends on the application. The 14-days free demo, downloadable from the Absoft website, can help to settle this question.

## Fx debugger

An important stage in the development cycle of a program is the

identification and the removal of possible bugs. A modern, fully featured debugger may dramatically speed up this process. Absoft offers *Fx*, which is one of the best debuggers we have tested so far.

According to the documentation, *Fx* is supposed to come with two interfaces, one *Motif*-based and the other *curses*-based. We wanted to try the latter, but apparently the executable wasn't part of the installation. This is a little annoyance that might signal some lack of accuracy in the documentation, but can be easily forgiven (although this thought possibly won't be shared by the few still insisting on just using a character console) when looking at the graphical incarnation of the debugger.

*Fx* supports C, *F77/90/95* and assembler, and features all the functionality you expect from a debugger, with as a bonus some nice additions like command logging and playback, software watchpoints and hardware breakpoints.

The main window (see screenshot, *above left*) has a neat appearance, with top menus enabling access to the functionality of the debugger and push buttons on the left for the most used commands, source code pane, output pane and a “run command” pane. From the main window a huge range of other features can be accessed, the review of which could easily fill an article by itself.

A product in its own right, *Fx* is a wonderful addition to *Pro Fortran*.

## Conclusion

*Pro Fortran 7.5* is a good product with excellent documentation and is pretty easy to install. The basic bundle features a well performing compiler and a marvellous debugger. A few glitches may be annoying, but do not overshadow the good job performed by Absoft. Recommended to those who need to go beyond *g77*. [LXF](#)

## LINUX Format VERDICT

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

A solid product with a good range of features. More than just a compiler.

## LINUX Format RATING

**8/10**

## LINK CHECKER

# LinkScan 10.0

Industrial strength link checking for webmasters. Neil Bothwick asks if it's worth the extra cost?

**Powerful, flexible link and site checking for medium to large sites. Competes with many free programs, but offers a lot more.**

- **PRICE** from \$750, see website for more details
- **PUBLISHER** Electronic Software Publishing Corporation (Elsop)
- **WEB** <http://www.elsop.com/>

**A**fter looking at the program details at the top of the page, your reaction was probably along the lines of "What! Five hundred quid for a link checker?", *LinkScan* will have to be good to justify any price tag, let alone one of \$750.

Installation worked well, with one exception. You have to run the installer as the *LinkScan* licensee. This worked until it tried to copy some scripts to the cgi-bin directory, which isn't writable by my normal user.

## Generating reports

Once installed, you run the *linkscan.pl* script and it churns through the pages in your **DocumentRoot**, checking links. You view the results as a series of reports in a web browser. Point it to the local CGI URLs you gave during installation and you'll get a menu of

the various wide ranging reports available. This is a very thorough program, both in its options and their docs. The Help link in the menu bar at the top of every report page is linked to the appropriate docs for that report.

Each report has an option to send it as an email; useful to keep a record of the reports, or forward error details to the person who must fix them. You can run *linkscan.cgi*, the program that generates the reports for the browser, as a shell process. This will write the report to a file or send it by email. This can be done interactively, by specifying the various options on the command line, or with a config file (a sample is supplied). A simple *cron* script could run *linkscan.pl* and then email various reports to your desk, to peruse while you drink the morning's first coffee.

I only found one real fault with the reports. I have the line **ImagePath = '/images/'**; within a **<script>** tag at the head of each page. *LinkScan* tried to follow this as a link and reported a 403 error as it's a directory not a file.

## Advanced config

The config created by the setup script will be sufficient for many users. To customise the operation of *LinkScan* you will need to edit the config files by

```

netz@marvin.digimed.co.uk: /var/www/linkscan - Shell - Konsole
Session Edit View Settings Help

Is LinkScan installed in a web-visible directory under www root? [y/n]: n

When prompted for a URL to the LinkScan directory, please supply the URL
to some other web-visible directory. You can create one later, if you prefer.
Then, copy the LinkScan docs/ directory (and the files therein) into that
directory so the LinkScan documentation can be viewed via a browser.
Enter the full URL. For example: http://www.example.com/linkscan/
Enter LinkScan URL: http://marvin.digimed.co.uk/linkscan

Does your server require that CGI scripts be installed in a special
directory (e.g. cgi-bin)?
[y/n]: y

Enter the *absolute pathname* of your special cgi-bin directory.
For example: /usr/www/cgi-bin/
Enter cgi-bin Pathname: /var/www/cgi-bin/

Enter the full URL to your special cgi-bin directory.
For example: http://www.example.com/cgi-bin/
Enter cgi-bin URL: http://digimed.co.uk/cgi-bin/

```

Initial configuration is done with a Perl script.

hand. There is a Preferences link in the menu bar, but this only covers a limited number of options. It would have been good to be able to alter all settings via the HTML interface.

*LinkScan* is certainly a useful and powerful program.. At this price, I would have liked to see a little more user-friendliness in installation and config, although the general operation is very straightforward. If you need a capable, configurable and easy to use means of testing a site, *LinkScan* is worth considering. **LXF**

## LINUX Format VERDICT

Ease of use	8/10
Features	9/10
Performance	9/10
Value for money	6/10

A powerful tool for webmasters, although installation and configuration could be easier at this price.

## LINUX Format RATING

7/10

## Available Reports

Your site in minute detail

There is a wide range of reports available, in varying levels of detail. Each error is flagged according to its importance, and you can alter the importance of specific errors in the configuration files.

<b>Project Summary</b>	Summary statistics for the current project
<b>Summary of All Projects</b>	Summary statistics for all configured projects
<b>Problem Documents</b>	List documents containing potential problems
<b>Selected Status Codes</b>	List errors of specific types
<b>Document Detail</b>	List all/selected documents
<b>All Pages Linking To ...</b>	Find pages that link to...
<b>Critical Errors</b>	List most critical errors
<b>Orphaned Files</b>	List orphaned files
<b>Detailed Errors</b>	List all/selected errors
<b>External History</b>	View history of an external link
<b>Changed Documents</b>	Compare two scans of the current project
<b>Redirections</b>	List a summary of redirections
<b>Search Documents</b>	Ad hoc searching: document-centric
<b>System Configuration</b>	Display current LinkScan configuration settings
<b>Search Links</b>	Ad hoc searching: link-centric
<b>LinkScan/QuickCheck</b>	View source code and detailed analysis of a document
<b>SiteMap</b>	Display LinkScan SiteMap

```

LinkScan/QuickCheck
Digital Media
Project: default
Description: Digital Media
Owner: All

Title: digital magazine - new technology you can't live without
Full URL: http://kosh.clara.co.uk/issues/issue02/contents.php
Document Size: 8858 Bytes
Last Modified:
Owned by: issues

00001 <!DOCTYPE HTML PUBLIC "-//W3C//DTD HTML 4.0 Transitional//EN">
00002 <html>
00003 <head>
00004 <title>digital magazine - new technology you can't live without</title>
00005 <meta http-equiv="Content-Type" content="text/html; charset=iso-8859-1">
00006
<meta name="keywords" content="digital, magazine, news, technology, new technology, mobile,
00007

```

Clicking on an error in the report shows where in the HTML it occurs.

# Herding Cats:

## A Primer for Programmers Who Lead Programmers

**Richard Drummond** examines a self-help guide for the programmer who has crossed the line between coding and management.

- **AUTHOR** J. Hank Rainwater
- **PUBLISHER** APress
- **ISBN** 1590590171
- **PRICE** £25.00

**G**ood programmers are often cruelly treated: they are promoted to management. Taken away from the job that they love – coding – they are left floundering in a new and alien world of budgets, deadlines and HR.

This book aims to help ease the culture shock of the programmer who has been punished in this way.

Right from the start the book takes a refreshing and no-nonsense look at management. It splits management into two distinct areas – leadership

and management – and helps you to develop both skills. The author quite rightly points out that the simplest way of improving a software team is to improve the manager. Hence, the book focuses on managing yourself just as much as managing your team.

This book is a practical guide. It teaches the new manager how to assess the capabilities and so make the best use of the talents of his team, and how to handle issues such as hiring and firing. It teaches how to organise and assign tasks, and how to handle meetings. To aid with time management, the author includes the source code to task manager application that he wrote to manage his own projects. Importantly, the book devotes a lot of space on technical



leadership of a project, highlighting the differences between the architecture and design of project and the importance of both.

Rainwater has 30 years experience in the industry and brings this to bear in a thought-provoking and witty book. The traps that unwary managers can fall into are amusingly illustrated with anecdotal extracts giving real world examples of these mistakes being made. The section classifying programmers into types according to their coding traits – the Architect, the Magician, the Slob and so on – in particular made me chuckle. This is a

valuable survival guide for the new and not so new manager alike, but I also think programmers in non-leadership roles would enjoy and profit from this book. If nothing else, it'll make let see things from the boss's point of view.

### Linux Format VERDICT

A witty and practical survival guide for those newly promoted to leading software projects.

LinuxFormat **RATING**  
 **9/10**

# Physics for Game Developers

If your latest game's spaceship moves with grace and realism of Pacman, perhaps you need this book, says **Richard Drummond**.

- **AUTHOR** David M. Bourg
- **PUBLISHER** O'Reilly
- **ISBN** 0-596-00006-5
- **PRICE** £28.50

**A**t last kids can be given a valid reason to study maths and physics at school – you can grow up to be a games programmer. After all, it's not enough that games look realistic, they have to feel realistic – and that depends on how accurately you model the physical systems involved. *Physics for Game Developers* aims to help those programmers inject a level of physical realism into their games.

The book deals tackles this problem

in two ways. It provides a crash course in mechanics to refresh your memory of the physics involved (it assumes a knowledge of calculus), and it shows how those equations can be translated into computer models. Example code is provided for *DirectX* on Windows, which should be translatable to other platforms. Basic topics covered include Newton's laws of motion, 2-D and 3-D particle kinematics, friction, collisions, and projectiles. The author takes the physics to a fairly advanced level, providing equations, explanations, and some guides to application. Later he uses all of this to build accurate computer models of the motions of aircraft, ships, cars and hovercraft.

The goal of this book is sound – to



bring together in one volume all the info a games programmer will need to create realistic games. It's a lot more digestible and practical than your standard physics textbook, but I don't think the author has struck quite the right balance. The treatment of the mechanics involved is necessarily brief. *e.g.*, the book glosses over numerical methods for differential equations, a crucial concept for any computer modelling. It would have been more beneficial to the reader to have more space devoted to the

actual task of building software models, rather trying and failing to cram in a derivation of all the mechanics involved.

### Linux Format VERDICT

A brave attempt, but doesn't hit the right balance between text book and cook book.

LinuxFormat **RATING**  
 **6/10**

# Rebel Code

Open Source – the history. **Richard Drummond** reviews a Penguin guide to Linux's roots.

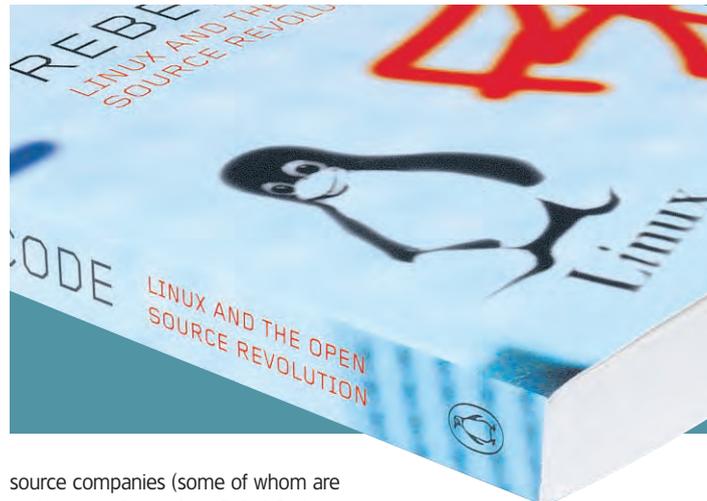
- **AUTHOR** Glyn Moody
- **PUBLISHER** Penguin Books
- **ISBN** 0140298045
- **PRICE** £7.99

**T**o truly understand the open source market today, you need to understand how it got here in the first place. And there's no better guide to the history of the open source phenomenon than Glyn Moody's *Rebel Code*, a book which charts the history of the open source movement throughout the 1990s. The book has been re-issued in a new paperback edition (rather aptly) by Penguin Books.

*Rebel Code* follows the open source movement from the humble beginning of such projects as the Linux kernel itself, *Apache* and *XFree86* to the dizzying heights reached by open

source companies like Red Hat and VA Linux – when stocks were soaring, and it seemed like the open source phenomenon could conquer the world. The first edition of the book was written before last year's Dot Com crash, and so the tone towards the end of the book inevitably had an almost religious zeal. This edition has been brought back down to earth with a new epilogue which gives a more pragmatic overview of the today's open source market and more realistic expectations of its future.

The author has based the book on interviews with the key protagonists in the open source world. Thus, the reader gets first-hand accounts from people like Linux Torvalds, Larry Wall and Eric Raymond. Accounts of the birth of all the major open source projects are featured, as well as the beginnings of some of the major open



source companies (some of whom are no longer with us, alas). *Rebel Code* is written with the lay person in mind, but is not patronising. It's lively and engaging style clearly reflects the author's passion for the subject matter. However, it doesn't fall into the trap of many popular science accounts: overdramatising the *personae* involved.

*Rebel Code* is an essential read for those new to open source, those who haven't actually lived through the history. More than just being an enjoyable and entertaining read, it will

help the reader understand the state of the open source market today.

## Linux Format VERDICT

An authoritative and entertaining history of the open source movement and now even better value for money in this new paperback edition

LinuxFormat RATING  
 9/10

# Java in a Nutshell, 4th Edition

Electronic docs are fine, says **Richard Drummond**, but a book is hell of a lot more portable.

- **AUTHOR** David Flanagan
- **PUBLISHER** O'Reilly
- **ISBN** 0-596-00283-1
- **PRICE** £28.50

**H**aving programming documentation in an searchable, electronic format is valuable, but sometimes there's just no substitute for hard copy. *Java in a Nutshell* fulfils this need, providing a comprehensive, 900-page source of reference to the Java language and the key Java APIs. The 4th edition has been updated to reflect the changes and new features in version 1.4 of the Java Platform, Standard Edition. This book covers the majority of J2SE 1.4's class library – including the New I/O, XML and Cryptography packages. What's not

included are the graphics and GUI packages: *AWT*, *Swing* and *Java 2D*. The first two have previously been covered by a sister-volume, *Java Foundation Classes in a Nutshell*, but that one has yet to be updated for J2SE1.4.

The book is divided into two parts: the first introduces the Java language's syntax and features and describes the basic usage of key APIs; the second part forms the main body of the book and contains the reference proper. The first nine chapters provide a guide to Java from first principles. It's not really intended as a tutorial, but it's great when you need to look up how some language feature works. Chapter 4 gives a quick useful overview of the Java class library, while the following two chapters cover security and JavaBeans, respectively. The final



chapter in the first section provides documentation for the standard Java development tools included with the SDK, such as *javac* and *jarsigner*.

The API reference is machine-generated and typeset directly from the Java APIs themselves, so accuracy should not be a problem. This section is ordered by package name, and each class is documented clearly with its fields, constructors and methods. With the 4th edition, however, ease of use has taken a fall: the page-edge marks, which previously let you easily flip to a

desired section, have disappeared. Hence, it now takes three takes as long to find the class you need. **LXF**

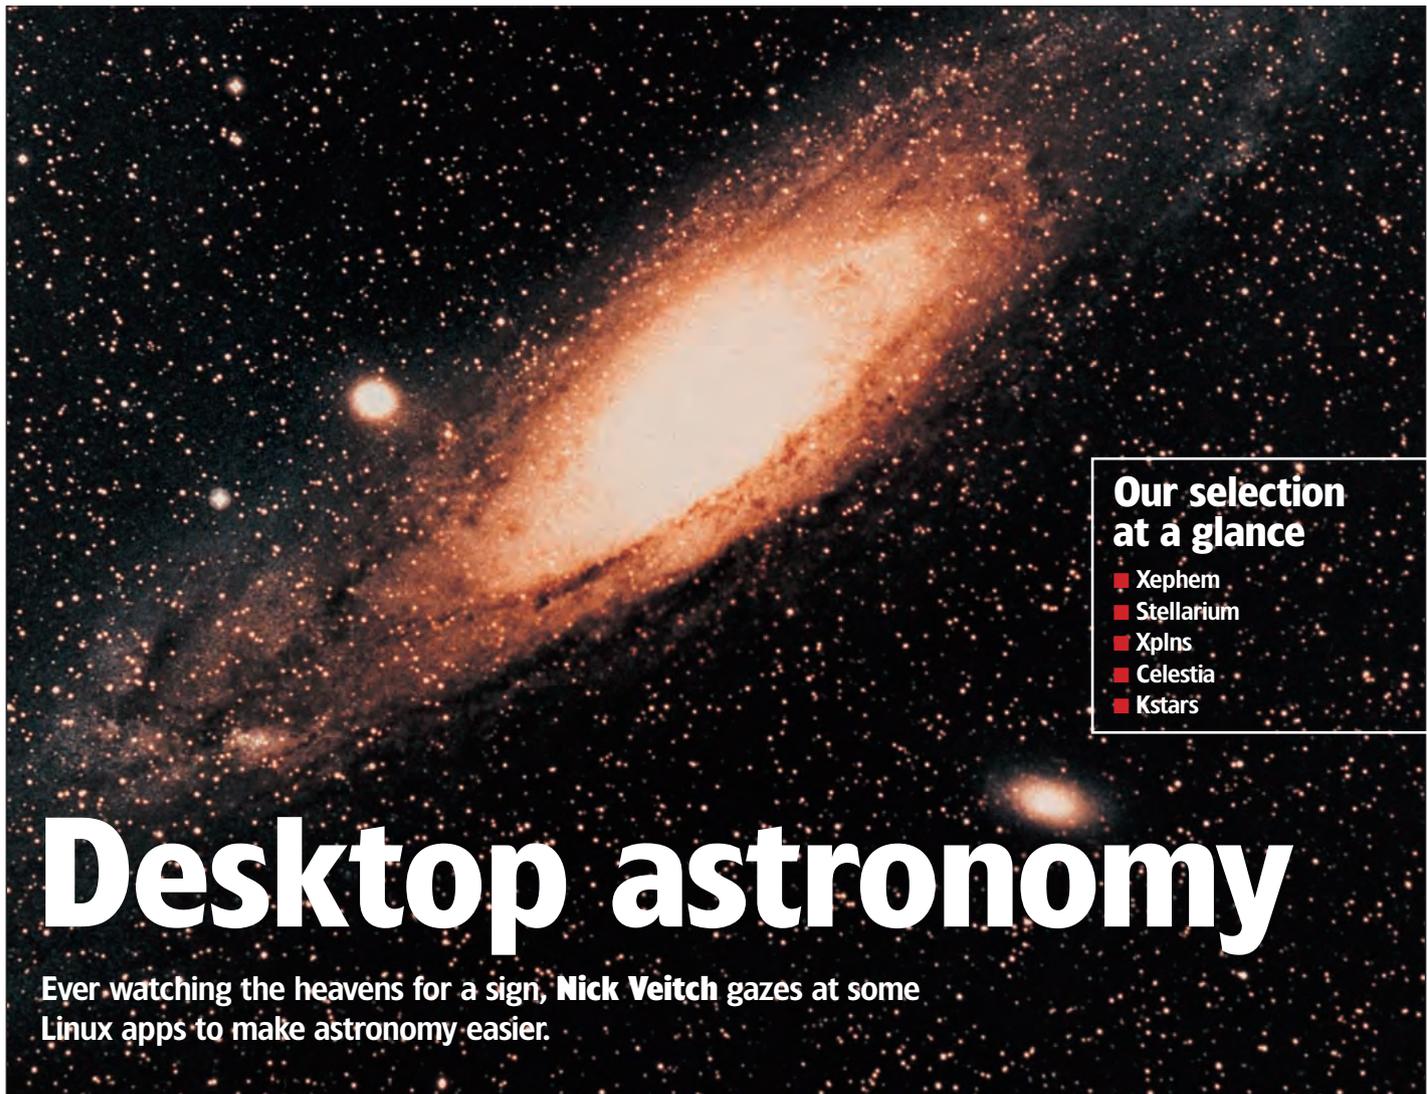
## Linux Format VERDICT

The navigation problems with the 4th edition are disappointing, but this is the only printed reference of its kind covering the Java platform and so is an invaluable book for Java programmers.

LinuxFormat RATING  
 8/10

# Roundup

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



## Our selection at a glance

- Xephem
- Stellarium
- Xplns
- Celestia
- Kstars

# Desktop astronomy

Ever watching the heavens for a sign, **Nick Veitch** gazes at some Linux apps to make astronomy easier.

Unix and Linux systems have always had a place in science. Unix servers have long been used by space agencies and academia for processing data of all sorts, and many of the tools they use are Free Software, some of which can trace their lineage back to the days when offline storage was a long strip of paper and a hole punch made a useful debugging tool.

For desktop astronomy needs though, we have to look a bit further than these accurate, but less immediate and user friendly, tools. The average hobbyist astronomer wants more than a command line tool

that will give the co-ordinates of a given star when supplied with an exact Julian date. Fortunately the last two years have shown an increase in projects of this type, so there is now a choice of desktop software to suit just about everybody.

But what do we want from such software? The ideal application should have a good range of star catalogues, include features such as naming of constellations, and provide information on every visible object, even if it's just an identification. A range of colour schemes, or a customisable palette, is desirable, because you may want to print out a chart without using

**“Some astronomy programs trace their lineage back to the days when a hole punch made a useful debugging tool.”**

up all the toner/ink. Also, for use at night, without destroying your night vision, you may want a red/black colour scheme.

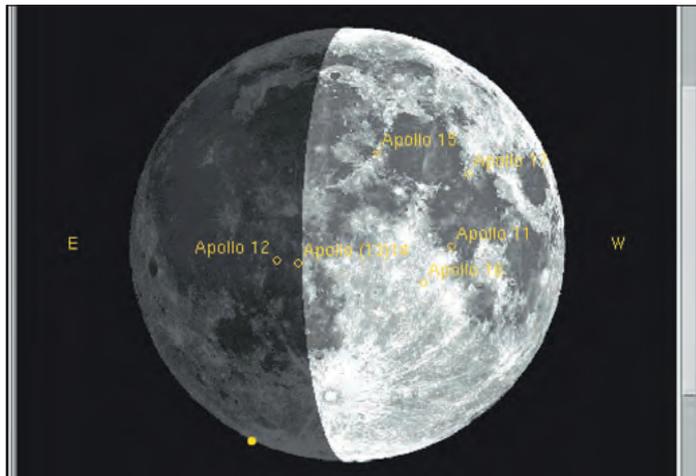
Speed in rendering the sky is desirable, though difficult to obtain. The main problem for software is that to draw the position of a star accurately requires several floating-

point trigonometric calculations. When you are dealing with thousands of stars, that can take some time. If you want to be really accurate you have to take into account all sorts of other anomalies, like the refractive nature of the atmosphere, etc. Often there can be a trade off between accuracy and speed.

# Xephem

As old as the stars

■ **VERSION** 3.5.2 ■ **WEB** [www.clearskyinstitute.com/xephem](http://www.clearskyinstitute.com/xephem)



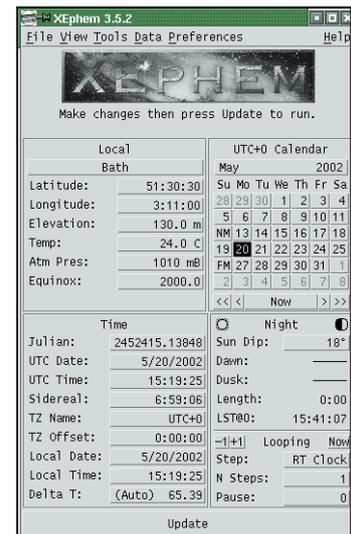
**Xephem provides separate views for the moon and major planets.**

**As it began its life in 1990,**

*Xephem* is probably the oldest desktop astronomy software on test here, so the developers have had plenty of time to get it right.

*Xephem's* view of the world is quite US-centric, so in the preprogrammed location data you'll find a hundred or so US cities, but few European ones. You can enter your own data though, and it includes information like the current temperature, barometric pressure and height above sea level, which will all have an effect on the exact positioning of the stars.

It's *Motif*-based interface may make it easy to create versions for other OSes, but isn't so pleasant to work with. It's also quite slow. Partly this is because of the high precision that *Xephem* works to. This gives you a really accurate view of the sky, and is very important when driving a telescope – which is another thing that *Xephem* has in its favour.



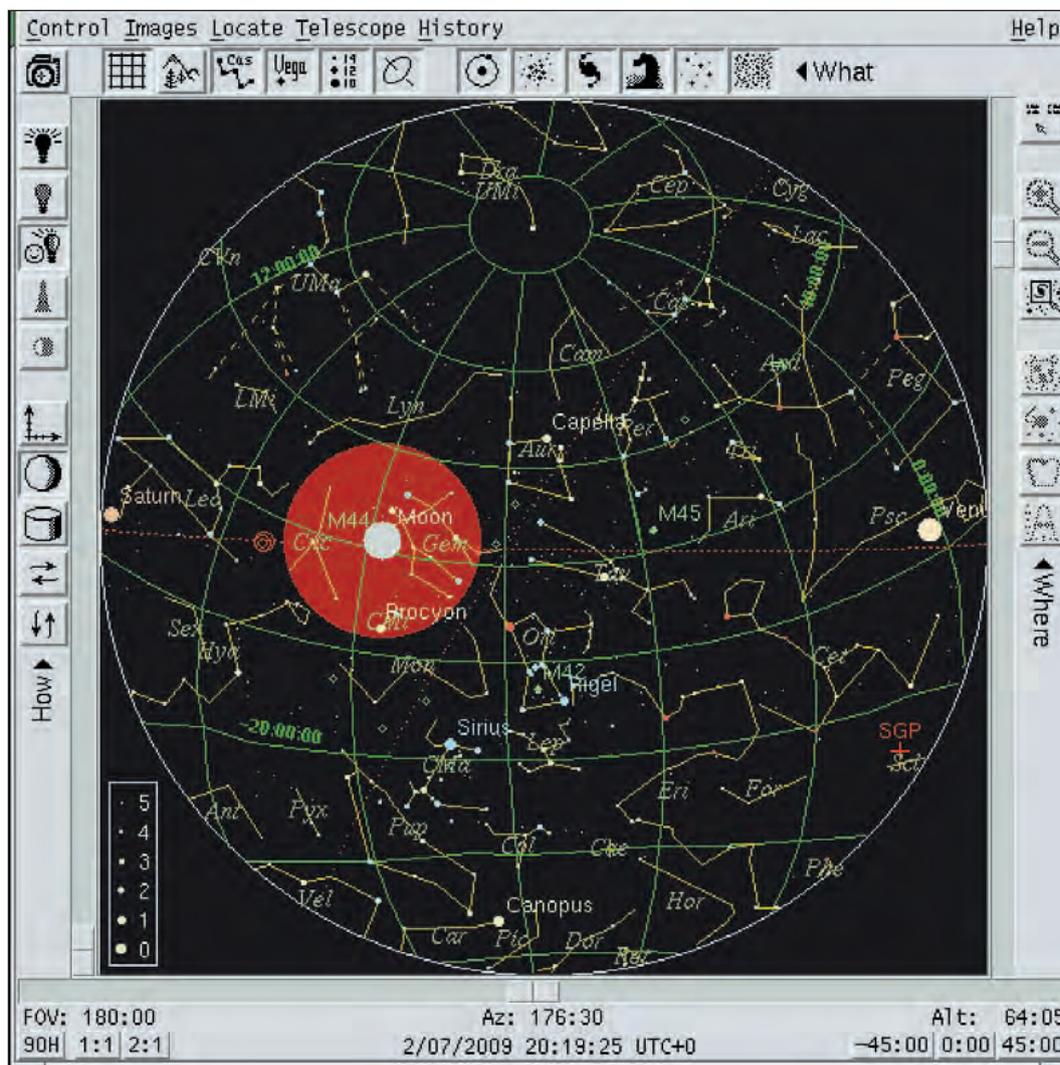
**the *Xephem* main panel isn't inspiring, but it is functional.**

Standard RS232 telescopes can be aimed with ease (we couldn't test this).

Panning around the sky is slow, even if you turn off the high precision mode – perhaps more selective control would help here. The Main sky view itself can be a little confusing at first, due to a lack of orientation devices (an Az/Alt readout is all you get). You can create a history of views though – more than most of the other packages allow – and it is possible to create your own "horizon" outline, to mark the location of landmarks (hills, houses etc), which will help you locate objects in the sky.

The views of solar system objects are good (there is a separate solar system mode, like an orrery), and there are Internet links for more data and pictures of objects.

One downer here is that *Xephem* is not truly "Open Source". Yes, you do get the source code, but its usage is restricted. *Xephem* probably has the most features, and the most accurate display, but it's slow and not quite as easy to use as other software featured.



**The red area indicates the view through a selected telescope eyepiece.**

## LINUX Format VERDICT

Installation	5/10
Documentation	7/10
Features	9/10
Ease of use	5/10

The best choice for the serious astronomer.

LINUX Format **RATING**  
**8/10**

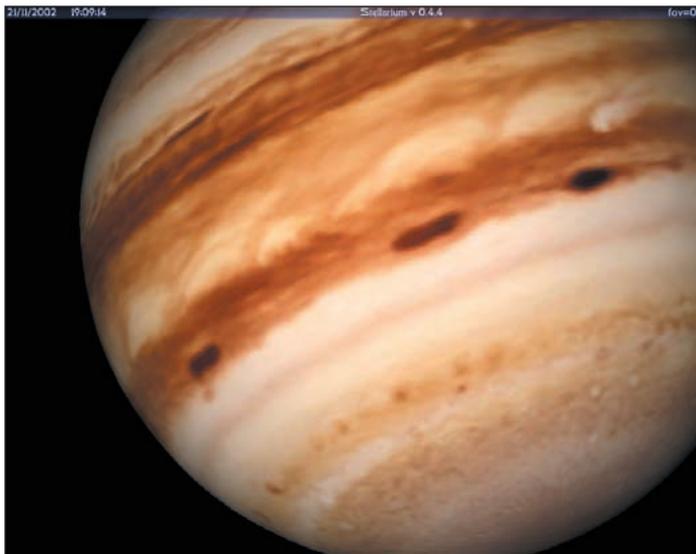
# Stellarium

## OpenGL based wonder

■ **VERSION** 0.4.4 ■ **WEB** <http://stellarium.free.fr/>

**This OpenGL-based software is designed for speed, not accuracy.**

However, if you are using astronomy software merely as a guide to where



Zoom right in to a planet and be treated to a texture mapped rendition.

things are in the sky, as opposed to trying to accurately direct a telescope, it does have some huge advantages. The star catalogues are rendered as individual GL objects, which means zooming and panning around the sky is lightning fast. If you aren't concerned with a very precise RA/Dec location for an object (and if you aren't using a high power telescope, you probably aren't) then the solution adopted here has a lot of advantages. The *Open GL* libraries take care of drawing the stars, and do so very quickly, taking advantage (hopefully) of the hardware acceleration of your graphics card. This means you can pan about the sky easily to locate objects, and zoom in and out with ease. It also means that visual effects such as fog, twinkling stars and graduated skies become possible.

Being completely *SDL/Open GL* though, all the controls are through the keyboard, or a manually edited config file. There aren't any comprehensive instructions either, which isn't so helpful, but an on-screen help menu will show you which buttons to press.

The display looks great. keyboard panning makes it a bit jerky, but in static modes, the twinkling stars (yes, they do twinkle), the graduated sky and simulated fog effect make for realistic viewing. An option to include your own 3D Horizon would have been a boon, for super-realism. Try the accelerated time modes for a look at how the sky changes over time!

There are lots of things missing though. No printing, no adjusting the palette, no searching for objects, and fairly limited info on the objects that are present. If you have a huge monitor and a fast graphics card though, this is an experience.

### LINUX Format VERDICT

Installation	5/10
Documentation	1/10
Features	7/10
Ease of use	8/10

Certainly the prettiest of the contenders, and pretty fast too.

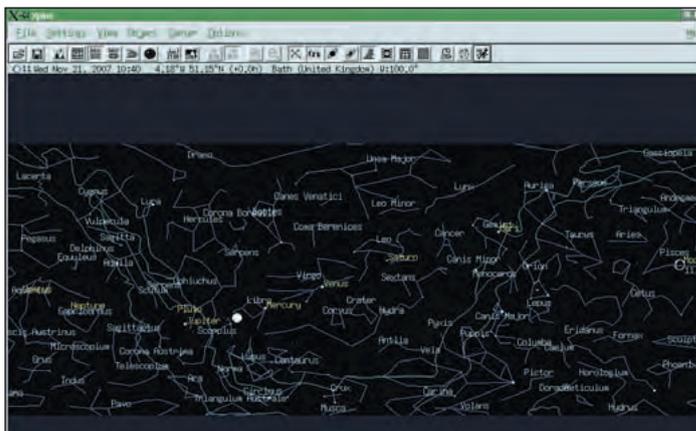
### LINUX Format RATING

6/10

# Xplns

## Lightning fast

■ **VERSION** 3.3.0 ■ **WEB** [www.astroarts.co.jp/products/xplns](http://www.astroarts.co.jp/products/xplns)



Xplns has a smooth animation mode – shame you can't see it in this pic!

**Developed in Japan by the same company responsible for *StellaNavigator* on the Windows platform, and other astronomy software, Xplns is free but, again, not open source. It is only distributed in binary format, as the authors wish to protect their astronomical libraries.**

This is a shame, but if you have an x86 machine, you can run the binary no problem. It isn't dynamically linked, so you don't need any libraries at all.

The first impressions of Xplns are wonderment that it is so fast. It seems to be even faster than the *Open GL* based software, while at

the same time being undoubtably more accurate.

The main interface is very clean (though the *Motif*-style widgets could do with a rethink), giving a generous view of the night sky. The software is missing any way of changing the colour palette, or printing out star charts, but at least you can save them as images. It is also possible to save 'Views' to come back to later – useful if you are planning for an astronomical outing in the future.

There is some 'click and drag' functionality to the main view, which is very fast, and doesn't blank out any of the objects when redrawing. Sliders on the right hand side and bottom control the view, but you'll have to resort to the keyboard or toolbar to zoom in and out. There doesn't seem to be any way of specifying the field of view, but the documentation isn't too good, so we may have missed it.

As well as the star view, there is a solar system view and a couple of "all-sky" views which are a bit weird.

It's also possible to add your own object catalogues, although the format doesn't seem to be explained anywhere. Right clicking on objects brings forth a wealth of information

and allows the user to centre, or lock onto, an object. The latter is useful when you are using the time controls. Xplns allows you to easily skip backwards and forwards in time (not literally, in terms of the sky display!) and you can even animate the display with different step times. The trace mode can then quickly show you what directions various objects are moving in.

On the minus side there is no easy way of changing the colour palette or any printing controls, so you are pretty much stuck to viewing the sky on screen unless you want to save images out and recolour them in *GIMP* or whatever.

### LINUX Format VERDICT

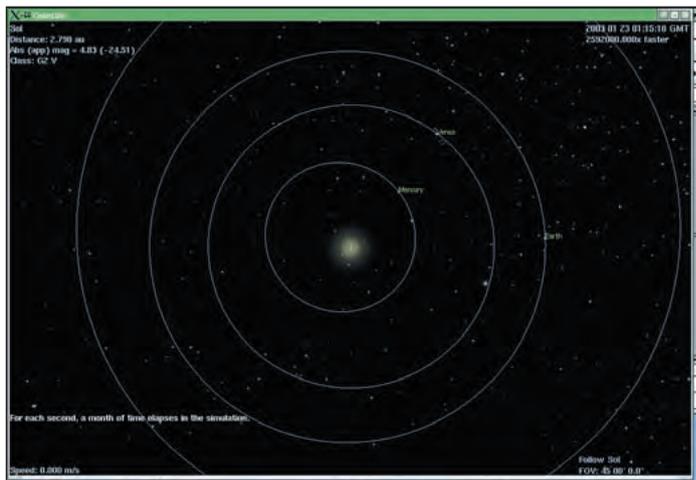
Installation	7/10
Documentation	3/10
Features	9/10
Ease of use	9/10

Excellent software, lacking in a few features.

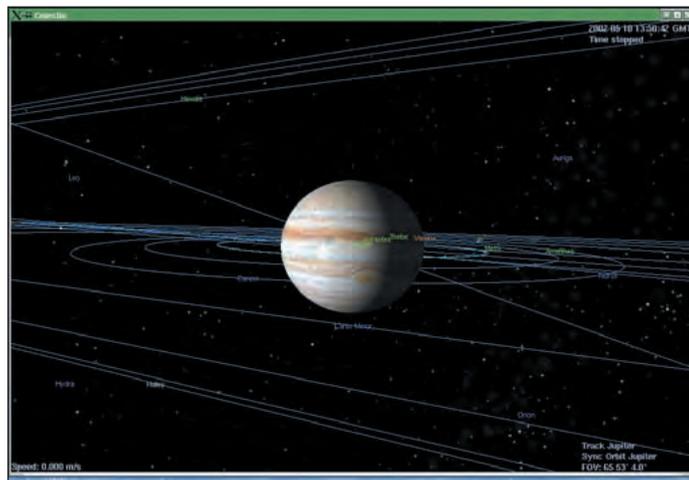
### LINUX Format RATING

9/10

# Roundup Astronomy



The demo mode takes you on a tour of the solar system.



The solar system objects are rendered in great detail.

## Celestia

### Universe simulation

■ **VERSION** 1.2.4 ■ **WEB** [sourceforge.net/projects/celestia](http://sourceforge.net/projects/celestia)

This is probably the odd-man out, more like the *openuniverse* software than the others featured here. An interactive 3D solar system explorer, you can also use it to conveniently map the stars in the night sky.

You find yourself at the controls of a free-flying camera, which you pilot

through the universe. A demo mode takes you to local areas of interest in a well made fly-by of the star system we call home. You can move off at warp speed through the universe and visit other bodies, like old favourites, Alpha Centuri and Vega. Controls are never adequately explained, but a quick

press of the keyboard should enlighten you. It's tricky to pilot the craft though, and quite easy to get 'lost in space', though it's surprising how you recognise distorted constellations.

*Celestia* will make use of both the Hiparchos and Tycho catalogues (which, space permitting, we have included on the CDs and DVD) so you can extend your universe.

Clicking on any celestial body will bring up useful info, but for some objects this info is a bit scant. Still, this is an entertaining and educational app, which may prove an interesting

companion to your mainstream night-sky software.

### LINUX Format VERDICT

Installation	6/10
Documentation	1/10
Features	7/10
Ease of use	6/10

An interesting project which has a lot going for it.

LINUX Format **RATING**  
 5/10

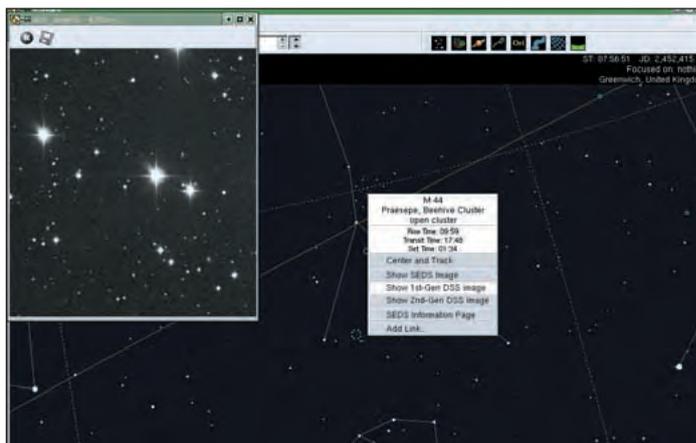
## Kstars

### The KDE option

■ **VERSION** 0.9 ■ **WEB** [www.kde.org](http://www.kde.org)

There are two versions of *Kstars* currently available: KDE2 and KDE3, though the latter contains no extra features over those offered by KDE3.

For different ways you might use the software, *Kstars* comes up trumps. There are several preset colour schemes and the ability to configure



*Kstars* links to DSS image sites and SEDS if you have Internet access.

more. The three defaults are suitable for general viewing, 'night-vision' and a chart scheme useful for printing out.

The main star display can be rotated using the mouse, or keyboard shortcuts will jump the display (useful for laptops). When positioning with the mouse, the stars are turned off, though you can still see lines and planets, otherwise you wouldn't know where you were panning to.

Right-clicking on an object brings up useful info like a name, type of object etc. A huge number of objects has weblinks to Digitized Sky Survey images at MAST, and other online resources. What is lacking is any other incidental information, even on planets (like radius, distance from earth etc). Some info is included in Internet links, but even so, some local info on the objects would probably be welcome.

Catalogue-wise, *Kstars* uses the fairly standard SAO data, with the Messier, NGC and IC objects. You can add further catalogues, providing you first convert them to a format *Kstars* will recognise. Tricky, but better than

not being able to edit the data at all!

One thing it lacks is the ability to 'save' a particular view. This would just mean saving the time and co-ordinate data, and be useful for reference.

For telescope users, the ability to define a field of view and magnification setting would be very useful. It's often difficult to match up what you are viewing on screen to what you see through a scope. Some software solves this by using a telescope cursor or template, which shows exactly what you should see.

### LINUX Format VERDICT

Installation	8/10
Documentation	6/10
Features	8/10
Ease of use	9/10

*Kstars* has come a long way in a short time, but still has some way to go to rival the best.

LINUX Format **RATING**  
 8/10

## Other projects worth watching

### Starplot

■ **WEB** [www.princeton.edu/~kmccarty/starplot.html](http://www.princeton.edu/~kmccarty/starplot.html)

#### The idea behind this little

application is to give you some idea of where you are in the grand universal scheme of things. In the middle, according to starplot.

It uses the Yale and Gliese star catalogues to create a three-dimensional view of space, which you can zoom in and out of and rotate, to get a better appreciation of the juxtaposition of stars in true

space (you might be surprised to find, for example that most of the stars in the 'Plough' are further away from each other than some of them are from Earth).

It's an interesting educational tool, but it wouldn't be wise to pilot your spacecraft using it. There is too little fine control of the zoom and rotation of the view, and labels appear and disappear when zooming, making it difficult to relate one view to the next. A little more control of what's displayed and how would make this a whole lot better, but it is an interesting project nevertheless.

### Marsproject

■ **WEB** <http://mars-sim.sourceforge.net/javadoc/MarsProject.html>

The *MarsProject* simulation software is a bit of an odd one. The premise is that colonisation of Mars has begun, at various bases on the Martian planet. Unlike Ray Bradbury's tales, there don't seem to be any Martian ghosts to interfere with the colony, so the simulated colonists go about their daily business, collecting samples, building things, growing food and maintaining their vehicles and biospheres. The

simulation keeps track of individuals and vehicles, and lets you know what they are up to. Each colonist has specific attributes (which you can edit if you want to put yourself on Mars, for example).

At the moment, that's about it. You can't direct the colonists to do anything, or interfere in any way, just observe.

Future versions may have more interactivity, but this is one of those projects you can easily get involved with, and help decide which things should be implemented next. Give it a try.

# ASTRONOMY THE VERDICT

Linux it seems, is well served with astronomy software, which ranges from the quick and easy general look at the night sky approach of *Stellarium* to the detailed precision *Xephem*.

For serious astronomy work, especially with a serial controlled telescope, *Xephem* is really the only choice. The interface may be ungainly, and parts of the program may be a little over-complicated, but it does seem to provide the most accurate results – imperative if you are to make sure you are looking at the right body in the heavens. Yes, it could be easier



**There is an orrery mode including a unique stereoscopic view function.**

to use, and possibly faster – if it had more options for deselecting accuracy features for example, allowing you to only use the most accurate modes when necessary. Currently there is only a Hi and Low accuracy setting.

If you want to have a quick look at the night sky to see what you can see, then *Stellarium* or *Xplns* are good options. *Stellarium* isn't as accurate, but it is very fast, if you have hardware acceleration, and can give you an idea of what the sky should look like on a particular evening.

*Xplns* on the other hand is very fast but doesn't need OpenGL. There are a few restrictions, but all in all the extra features (like the automatically generated 'Things of Interest' table) make this a more rounded application.

*Kstars* has a lot to recommend it – you can change the palette, add your own star catalogues, and easily print out the resulting images. It isn't as fast as the others, but shows promise, being more actively developed than the likes of *Xplns*, and it is truly open source, which is always a bonus.

In summary, *Xephem* is the best for accuracy, and essential if you have a computer-controlled telescope. Of the others, *Xplns* is probably the best, but *Stellarium* is faster. *Kstars* is worth keeping an eye on. **LXF**

### Table of features

Name	Celestia	Stellarium	Kstars	Xephem	Xplns
Version	1.2.4	0.4.4	0.9	3.5.2	3.3.0
GUI	OpenGL	OpenGL	KDE/Qt	Motif	Motif-like
OpenGL rendering?	Yes	Yes	No	No	No
Planets	Yes	Yes	Yes	Yes	Yes
Orrery mode?	Yes	No	No	Yes	Yes
Saved Views	No	No	No	History	Yes
Nightview mode	No	No	Yes	No	No
Chart printing	No	No	Yes	Yes	No
Twinkling?	Yes	Yes	No	No	No
Timelapse mode	No	Yes	Yes	Yes	Yes
Object searching	Limited	No	Yes	Limited	Yes
Object tracking	Limited	Limited	Yes	Limited	Yes
Web links	No	No	Yes	Yes	No
Telescope control	No	No	No	Yes	No
Additional catalogues?	No	No	Yes	Yes	Yes
Adjustable horizon	No	No	No	Yes	No
<b>Overall Rating</b>	<b>5</b>	<b>6</b>	<b>8</b>	<b>8</b>	<b>9</b>

# HotPicks

The best new open source software on the planet!



**Richard Drummond**

Our reviews ed puts in a plea for Java and the framebuffer.

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 coverdiscs, 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](mailto:linuxformat@futurenet.co.uk)

## HotPicks at a glance

Jmol	38
XBindKeys	39
Epiphany	39
jEdit	40
XCDRoast	41
KSensors	41
XDirectFB	42

## HotPicks award

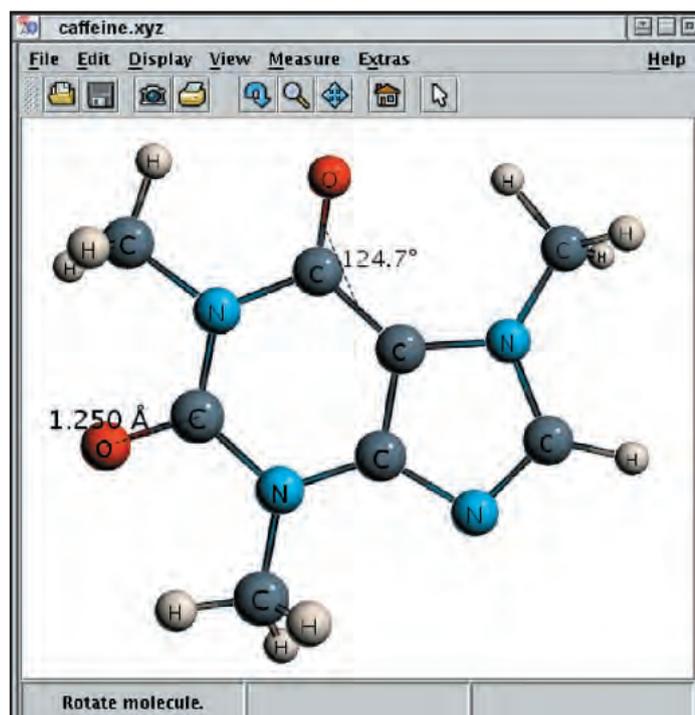
Everything covered in our Hot Picks section is unmissable, but every month we'll be singling out one project for outstanding brilliance. Only the very best will be chosen!



## CHEMICAL MOLECULE VIEWER

# Jmol

■ VERSION 3 ■ WEB <http://jmol.sourceforge.net/>



Jmol displaying one of my all-time favourite molecules, Caffeine.

Students of chemistry will probably be familiar with those kits with coloured balls and springs, used to build visual models of molecules. I'm sure you know the sort of thing I mean. They're a right pain, aren't they? Ten minutes into building your lovingly-constructed Buckminster Fullerene ball, you find don't have enough carbon atoms to complete it. You then have two choices: either steal some from your neighbour or adopt a virtual approach with a program such as *Jmol*.

*Jmol* is a pure Java application for visualising chemical molecules. The key word here is 'visualising': *Jmol* doesn't let you build models of molecules, it lets you view them in 3-D. It can, however, import molecule descriptions created in programs such as *XMol*, *Ghemical* and *Jaguar*. Lots of

sample files are included to demonstrate *Jmol's* capabilities.

*Jmol* is available as a standalone Java application or – in a slightly cut-down form – as an applet. The latter is great for teaching purposes. Imagine an online chemical text-book, where instead of having flat, static images of molecules – you have interactive, 3-D models of molecules which the reader can arbitrarily zoom and rotate. Both versions require a few extra Java class libraries to work, and the standalone version needs *Swing*, but these are all included with the package.

The main display of *Jmol* is the 3-D view of your chosen molecule. It doesn't use *OpenGL*, but nevertheless renders quickly. Various rendering styles are supported including wireframe, flat and shaded – and anti-aliasing can be turned on or off – so

you can tweak the display according to your CPU power and JVM capabilities. Atoms can optionally be labelled with their symbols, atomic numbers or charges. The molecule itself can be rotated, zoomed and translated by dragging with the mouse. The default mode is rotation, and the applet lets you zoom and translate by holding down the **Shift** and **Ctrl** keys, respectively. In the application proper, however, the mode may only be changed via the toolbar. Molecules may also be animated, for example, to simulate vibration of a molecule. In addition, *Jmol* includes various tools to calculate inter-atomic distances, bond angles and dihedral angles, and these can be calculated and displayed on the molecule and updated as the molecule is animated.

*Jmol* has good support for exporting images. You can save a rendered molecule as GIF, JPEG, PPM, BMAP or PNG image, print it or export it as *PostScript* file. One missing feature is the ability to save a molecule in a readily usable vector format, such as EPS or SVG; this would make using *Jmol's* output in your documents easier. *Jmol* is also able to export a molecule as a *POV* script, which can be then be used by the *POV* raytracer to render a high quality 3-D image of the molecule. This function currently has some teething problems, however. For instance, I could only get it to fly with the latest beta version of *POV* (version 3.5) – although which version it is supposed to support was not mentioned anywhere in the docs – and this worked only after I corrected the syntax errors in the generated script. The results look great, however, and consequently this is an area of the program I would like to see extended. The *POV* export dialog should offer more options to control the output, such as choosing the size and quality of the rendered image.

Despite the faults in the *POV* export function, *Jmol* is an easy-to-use application, which is sure to be a great time-saver for both students and teachers of chemistry.

## HOTKEY DAEMON

# XBindKeys

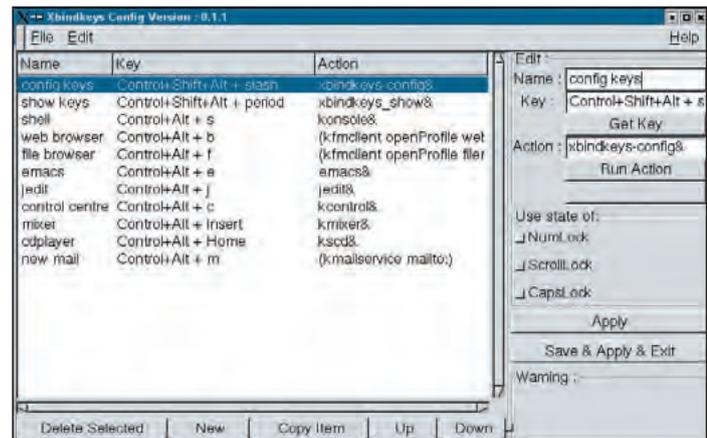
■ VERSION 1.5.5 ■ WEB <http://hocwp.free.fr/xbindkeys/>

**B**efore I was a Linux enthusiast, my operating system of choice was AmigaOS. Although proprietary, it had much to recommend it, and, in the early nineties was streets ahead of the competition on the desktop.

One of the Amiga's many unique features was a system called commodities – a mechanism for implementing desktop daemons with a consistent API and user interface (much like tray applets in today's desktops). The commodities system was used for many purposes – from screen blankers to window focus policy tools – but one popular use was the *hotkey* daemon, a program that would globally assign user-definable functions to key-stroke combinations.

The ability to launch arbitrary commands quickly from the keyboard without having to reach for the mouse is such a labour-saving device that I often wished there were some way to replicate this functionality under X. Well, it turns out that there is, with a program called *XBindKeys*.

*XBindKeys* is a small shell command which reads its configuration from a text file and sits in the background examining X input events. It is simple to set up: you just need to specify the shell command to be launched for a particular key and mouse button combination. To help you work out what labels X gives the various keys on your keyboard, you can run *XBindKeys* in a non-daemon mode to capture, identify and display key presses. To make things even



**XBindKeys Config** make light work of assigning hot keys to commands.

easier, a separate *GTK*-based configuration GUI is available, called *XBindKeys Config* (see [www.netchampagne.com/xbindkeys\\_config](http://www.netchampagne.com/xbindkeys_config)). This is none too pretty, nor terribly intuitive, but certainly speeds up the task of configuring your set-up.

*XBindKeys* is an immensely useful and simple tool. While there are various other tools that perform some of *XBindKeys*'s functionality – for

example, see *KHotkeys* (<http://dforce.sh.cvut.cz/~seli/en/khotkeys/>) – the advantage of *XBindKeys* is its generality: it is not tied to any particular desktop environment. Just start the daemon from your *xinit* script, and the same key definitions can be used under KDE, GNOME or *WMaker* sessions. A more polished configuration GUI would be a valuable addition, though.

## RETRO GAMING CLASSIC

# Epiphany

■ VERSION 0.3.0 ■ WEB <http://epiphany.sourceforge.net/>

**W**hoever said they don't make them like they used to was wrong when it comes to open source games: they patently do. Dozens of classic games have been revisited as open source projects, and one such newly established project – a clone of that old 8-bit favourite, *Boulderdash* – is enigmatically called *Epiphany*.

The aim of *Boulderdash* – and hence *Epiphany* – is to run around a series of mazes collecting all the jewels, avoiding any falling boulders and safely making it to the exit each time. In *Epiphany* there's no time limit, but you do have to collect the correct number of jewels before you are allowed to pass through each level's exit. Doesn't that sound simple? Did I mention that there are exploding bombs, locked doors and wacky little aliens, called *Peperons*, to overcome too?

It's been done before, yes, but *Epiphany* is a polished – if not exactly feature-packed – version of *Boulderdash*. Who minds a repeat when it's enjoyable fun?

It currently lacks any sound, but the graphics – especially since the latest release – are exceptionally cute and colourful. The level design is cunning, requiring just the right mix of strategy and digital dexterity (level five in particular is evil!).

There are currently only ten levels supplied with the game and no level editor, but, since the maps are plain text files, it shouldn't prove too much trouble to craft your own if you finish these.

*Epiphany* is elegantly coded in C++ and based on the *ClanLib* cross-platform gaming library. It has no other significant dependencies, so building it yourself is straightforward. Debian packages are also available for download.



**Epiphany** is polished version of that timeless classic, *Boulderdash*.

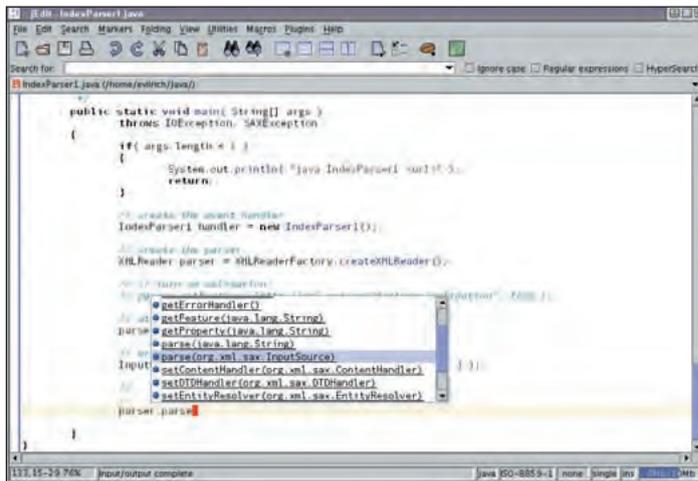


**That falling red blob is a bomb! I'd better scarp before it goes off.**

## PROGRAMMER'S TEXT EDITOR

## jEdit

■ VERSION 4.0 ■ WEB <http://www.jedit.org/>



The *SpeedJava* plug-in, offering powerful code-completion for the Java programmer, is one of many plug-ins available.

At the risk of being branded a heretic, I am going to suggest that *jEdit* is currently the best text editor available for Linux. Shock, horror! Yes, it's better than *Emacs* (although *Emacs* will forever have a place in my heart), and it's better than *Vi*. What makes this statement even more controversial is that *jEdit* is implemented in 100% pure Java. Have I totally lost my wits? Get over your preconceptions concerning Java on the desktop. *jEdit* really is that good.

So, what's so great about *jEdit*? Well, it's fast, it's configurable, it's visually appealing and it has all the

functionality you'll ever need – and then some. *jEdit* is extensible via macros written in BeanShell, an implementation of the Java language designed for scripting and executed in a built-in interpreter, and via plug-ins written in Java. Of course, you don't need to be a programmer to use macros, since a macro record facility is included. (A Jython interpreter is also available as a plug-in if you prefer macros written in a flavour of Python). A BeanShell console is provided for testing scripts and entering quick one-off routines. *jEdit* hasn't quite reached the level of *Emacs* in terms of overloaded



functionality, but a file-manager is included – which can double as a FTP browser with an appropriate plug-in. Other interesting plug-ins include an IRC client and an accounting ledger!

The editor proper features unlimited undo/redo, syntax-highlighting for dozens of languages, auto-indentation, folding, bracket matching, incremental and regular-expression-based searching – the list goes on. What's more, it's very easy to tweak *jEdit* to create a comfortable editing environment. Nearly every aspect of the editor can be configured – from colours to cursor style and from keyboard shortcuts to the toolbar. *jEdit* does not have a set of built-in personalities to mimic the feel of other popular editors like some others offer, but this is a minor complaint; you can always reconfigure it yourself to feel more like *Emacs* or *Brief* or whatever.

Multiple files can be edited at once, and *jEdit* will choose and work in an appropriate mode for each. Buffers can be displayed in separate windows (*jEdit* calls these 'views') – or you can split a single view either horizontally or vertically to work on more than one file (or view multiple parts of the same file) in one view.

Unusually, for such a practical tool, *jEdit* looks pretty. It has a well-designed, comfortably laid-out interface. Support for themes is provided by a plug-in. The text buffer can anti-alias fonts, so your code looks good, too, and *jEdit* produces beautiful hard copy. Printouts can include syntax-highlighting and (optionally) headers, footers and line numbers.

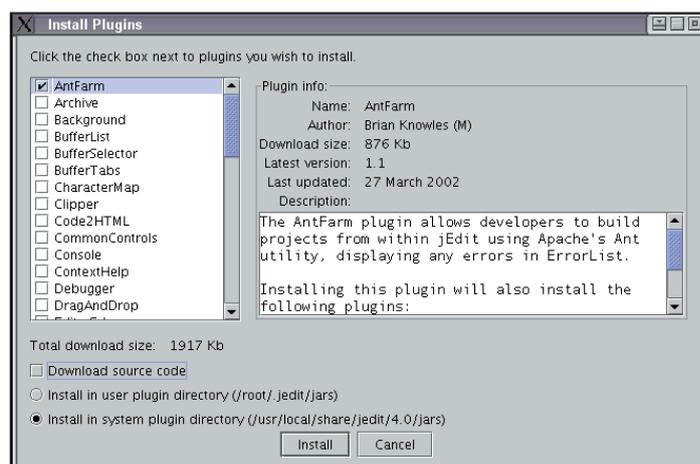
The *jEdit* community is particularly active, and the community site (<http://community.jedit.org>) hosts

downloads, mailing lists and other resources. You can find scores of macros and plug-ins ready to fetch, install and use. Official plug-ins are available from <http://plugins.jedit.org>.

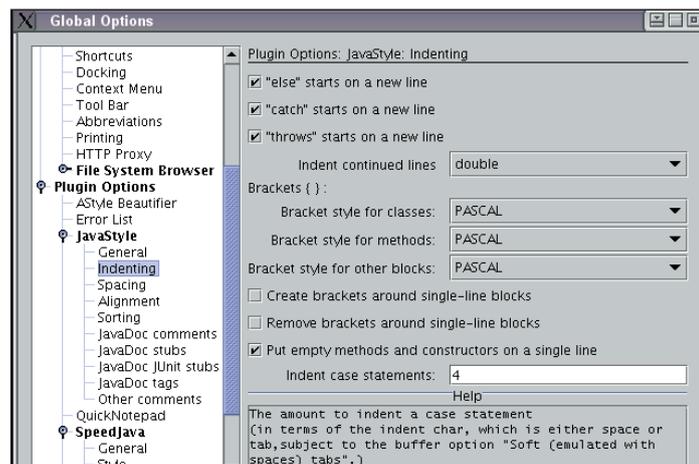
*jEdit* can check online here for new and updated plug-ins and automatically download and install them for you. Current plug-ins include spell-checking, templates, project management and XML tools.

The area that is best supported with plug-ins is – without doubt – Java development. For example, the *AntFarm* plug-in integrates with *Ant* for project building, the *SpeedJava* plug-in provides code-completion facilities, *JavaStyle* provide Java source code beautification, the *JIndex* plug-in provides tools to look-up and browse the Java API docs, and *JSwatPlugin* integrates with the *JSwat* Java debugger. Other languages are less well represented with tools, but, for example, the *Class Wizard* can generate skeleton code for both Java and PHP classes, while the *Tags* plug-in provides tools for working with tag files. My only complaint is that nobody has yet implemented *Tetris* or *Adventure* as a *jEdit* macro.

*jEdit* is great all-round text editor. It has a shallow learning curve, and this is aided by copious online docs covering everything from the basics to macro and plug-in writing. For the Java programmer, it can form the hub of a powerful development environment, offering integration with tools such as *Ant* and *JSwat*. Yes, *jEdit* is written in Java – and it won't really be happy on a machine with less than 128MB of memory – but as long as you have enough RAM, execution speed is not issue with a modern Java run-time.



Downloading and installing new plug-ins is a breeze with the built-in manager. You'll soon be getting productive.



*jEdit* is infinitely configurable, but the hierarchical Options dialog makes this complexity manageable.

## CD WRITING APPLICATION

# XCDRoast

■ VERSION 0.98alpha10 ■ WEB <http://www.xcdrtoast.org/>

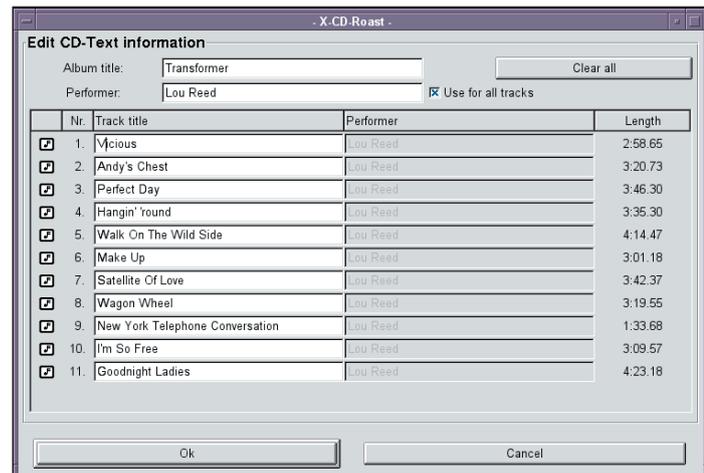
When we did the round up of CD burning software in LXF 25, we pointed out several shortcomings in the then current version of *XCDRoast*, one of the most popular CD writings apps for Linux. Many of these have now been rectified with this new release, version 0.98 alpha10. You'll notice it still has the worrisome *alpha* tag, but again has proved as solid as a rock under testing.

*XCDRoast* now has full support for writing multi-session discs, which is perhaps the most important feature that was missing from previous versions. Also, it now has support for reading, editing and writing CD-Text. This means you can get artist and track title listings from the web via a CDDDB server or enter it manually

yourself and write this information to disc when creating audio CDs, ready for playback in compatible devices.

Other tweaks include a new over-burning function and experimental support for writing DVDs.

The *XCDRoast* GUI has had modifications in this release which make it slightly easier to use. For instance, you now no longer have to remember to tell *XCDRoast* to calculate a track's size when mastering a data track: there's an option to do it automatically. Similarly, you don't have to explicitly accept your track layout on the 'Layout tracks' pane of the 'Write Tracks' page: switching to the 'Write tracks' pane implicitly accepts your layout. The configuration of *XCDRoast* for use by non-root users has been



***XCDRoast* is catching with its rivals, now supporting multi-sessions writings and CD-Text.**

made much simpler in this version, too. There's still work to be done, however. For example, leaving the 'Write tracks' page still makes *XCDRoast* – infuriatingly – forget your track layout.

*XCDRoast* has caught up in leaps and bounds with the competition with

this release. It still lags behind rivals such as *GnomeToaster* in tools for handling audio tracks – significantly, you still cannot use audio files other than CDDA-frequency WAVs as the audio source – but *XCDRoast* is certainly one of the most capable apps for mastering data tracks.

## HARDWARE MONITORING

# KSensors

■ VERSION 0.6.1 ■ WEB <http://ksensors.sourceforge.net/>

*KSensors* is a graphical front-end to the *LM Sensors* project for the KDE desktop. The *Sensors* project (see [www2.lm-sensors.nu/~lm78/](http://www2.lm-sensors.nu/~lm78/)) provides kernel drivers for a whole host of

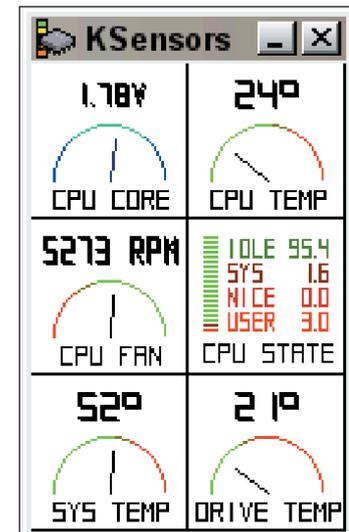
hardware sensors, such as any hardware monitoring functionality built into your motherboard, and allows you to monitor such properties (depending on your hardware) as CPU temperature,

CPU voltage and fan rotation speed. The idea is that keeping track of parameters like these will allow the early detection of hardware failure.

*KSensors* builds under both KDE2.2 and KDE3.0, and provides a tray applet that lives on the *Kicker* panel and a main window that can show the output from various types as sensors as temperature gauges, voltage gauges and so on. It also supports the *hdtemp* program (see <http://coredump.free.fr/linux/harddrive.html>) which can monitor the temperature of supported IDE hard drives by reading SMART information (only the more modern drives have temperature sensors, though), and can display various other system variables such as memory usage and CPU load.

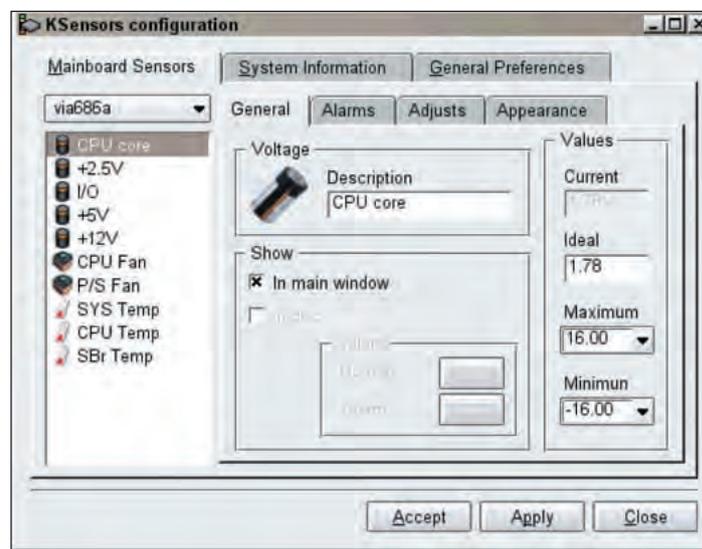
Setting up *KSensors* is a simple matter with its configuration dialog, and this gives the user a lot of power to control its appearance and behaviour. Each monitor can have an alarm associated with it, which, when the monitored value falls out of the permitted bounds, can either play an alert sound or run a specified program.

*KSensors* is a well-implemented and neat little tool, and is just perfect for those that want lots of gadgets,

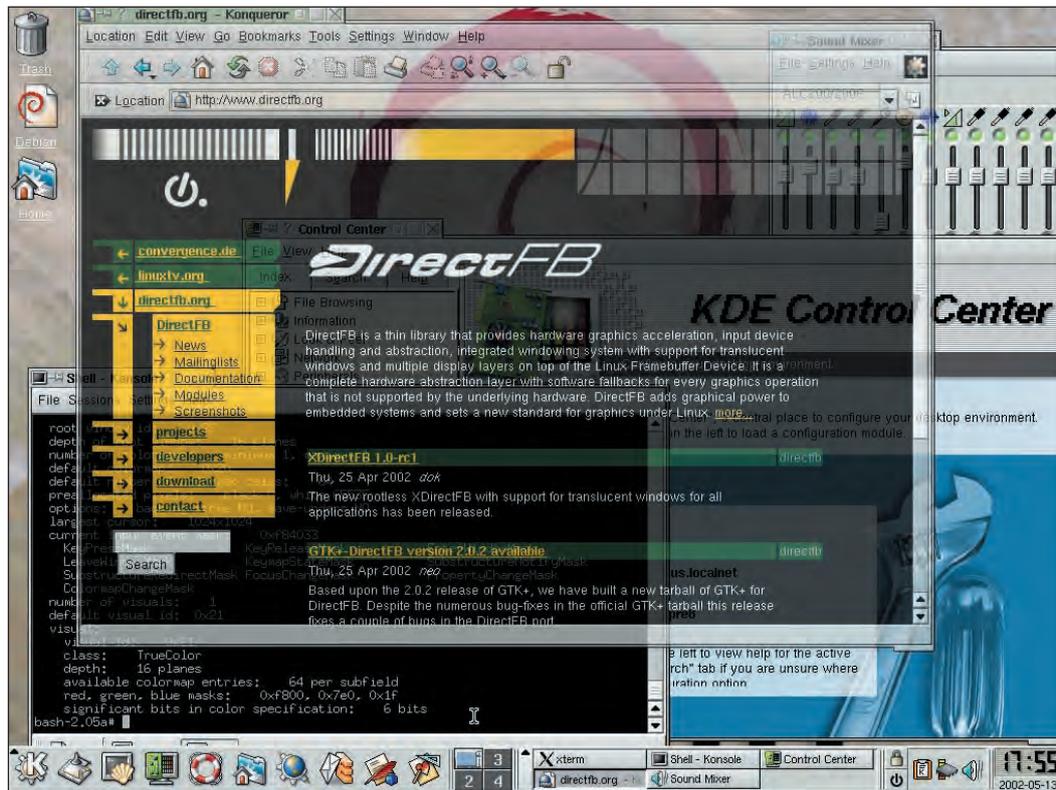


***KSensors* can display some important-looking gauges on your desktop – whether it has much practical value is up to you.**

gizmos and gauges on screen. The sensor daemon included in the *LM Sensors* package is a lot more practical, especially for servers – it runs in the background, logs sensors readings and can email you when an alert occurs. *KSensors*, however, is more suited to the desktop user (especially those pushing the limits of their hardware), and is attractive and easy-to-use way of keeping track of your hardware's status.



The configuration dialog gives you lots of scope to tweak *KSensors'* look and behaviour.



***XDirectFB* may not yet have much practical value, but those translucent windows look cool.**

## ROOTLESS X SERVER

# XDirectFB

■ **VERSION 1.0RC1** ■ **WEB** [www.directfb.org/xdirectfb.xml](http://www.directfb.org/xdirectfb.xml)

The *DirectFB* project is a modern, light-weight graphical library built on top of the Linux kernel's framebuffer layer – providing hardware-accelerated drivers for graphical operations with pixmaps (blitting, blending, stretching and so on), rather than bitwise raster ops. Consequently, it is particularly good at multimedia apps – and in fact is designed with digital TV in mind.

*DirectFB* would also make a solid basis for a multimedia desktop. It's not ready for this yet – it needs more polish and broader driver support – but something else it also needs is applications. Porting apps from X to *DirectFB* would be time-consuming and tedious, so the developers of *DirectFB* have addressed this problem in two ways: a port of *GTK+ 2.0* that uses *DirectFB* as its hardware layer – allowing any *GTK+ 2.0* applications to be recompiled for *DirectFB*; and a port of *XFree86* that runs on *DirectFB*, *XDirectFB*. With the latter, you don't even need to recompile apps – any X client will run on *XDirectFB* (providing

that it doesn't use any exotic extensions to the X protocol).

*XDirectFB* is a rootless X server. This means that X doesn't take over the whole display like it usually does. Instead, under *XDirectFB*, X windows are mapped onto native *DirectFB* windows – so you can run applications of both types on the same screen. This idea will be familiar to the users of Darwin port of *XFree86* for MacOS X, another rootless X server which lets you run X clients and native OS X apps together. *XDirectFB* also supports a root-window mode, so you can also use it like a conventional X server.

What can you do with *XDirectFB*? It's fairly complete, and supports a good portion of the extensions to the X protocol that are commonly used, so most software will run quite happily. The latest version even handles shaped windows. Notably, at the moment, there's no support for the **X Render** protocol (so no anti-aliased fonts), no **XVideo** (so no video capture or hardware-assisted playback of videos), no **XVidModeExtension** (so

no changing of screen modes), no **GLX** (so no *OpenGL* rendering, of either the software or hardware based variety) and no **DGA** (so games cannot directly access the frame buffer, e.g. for faster rendering). Having said that, all these types of features (with the exception of hardware-accelerated GL) are implemented right now in *DirectFB* itself. In theory, it shouldn't be too difficult to support them in *XDirectFB*; in the meantime, you'll have to port to *DirectFB* to use these kinds of features.

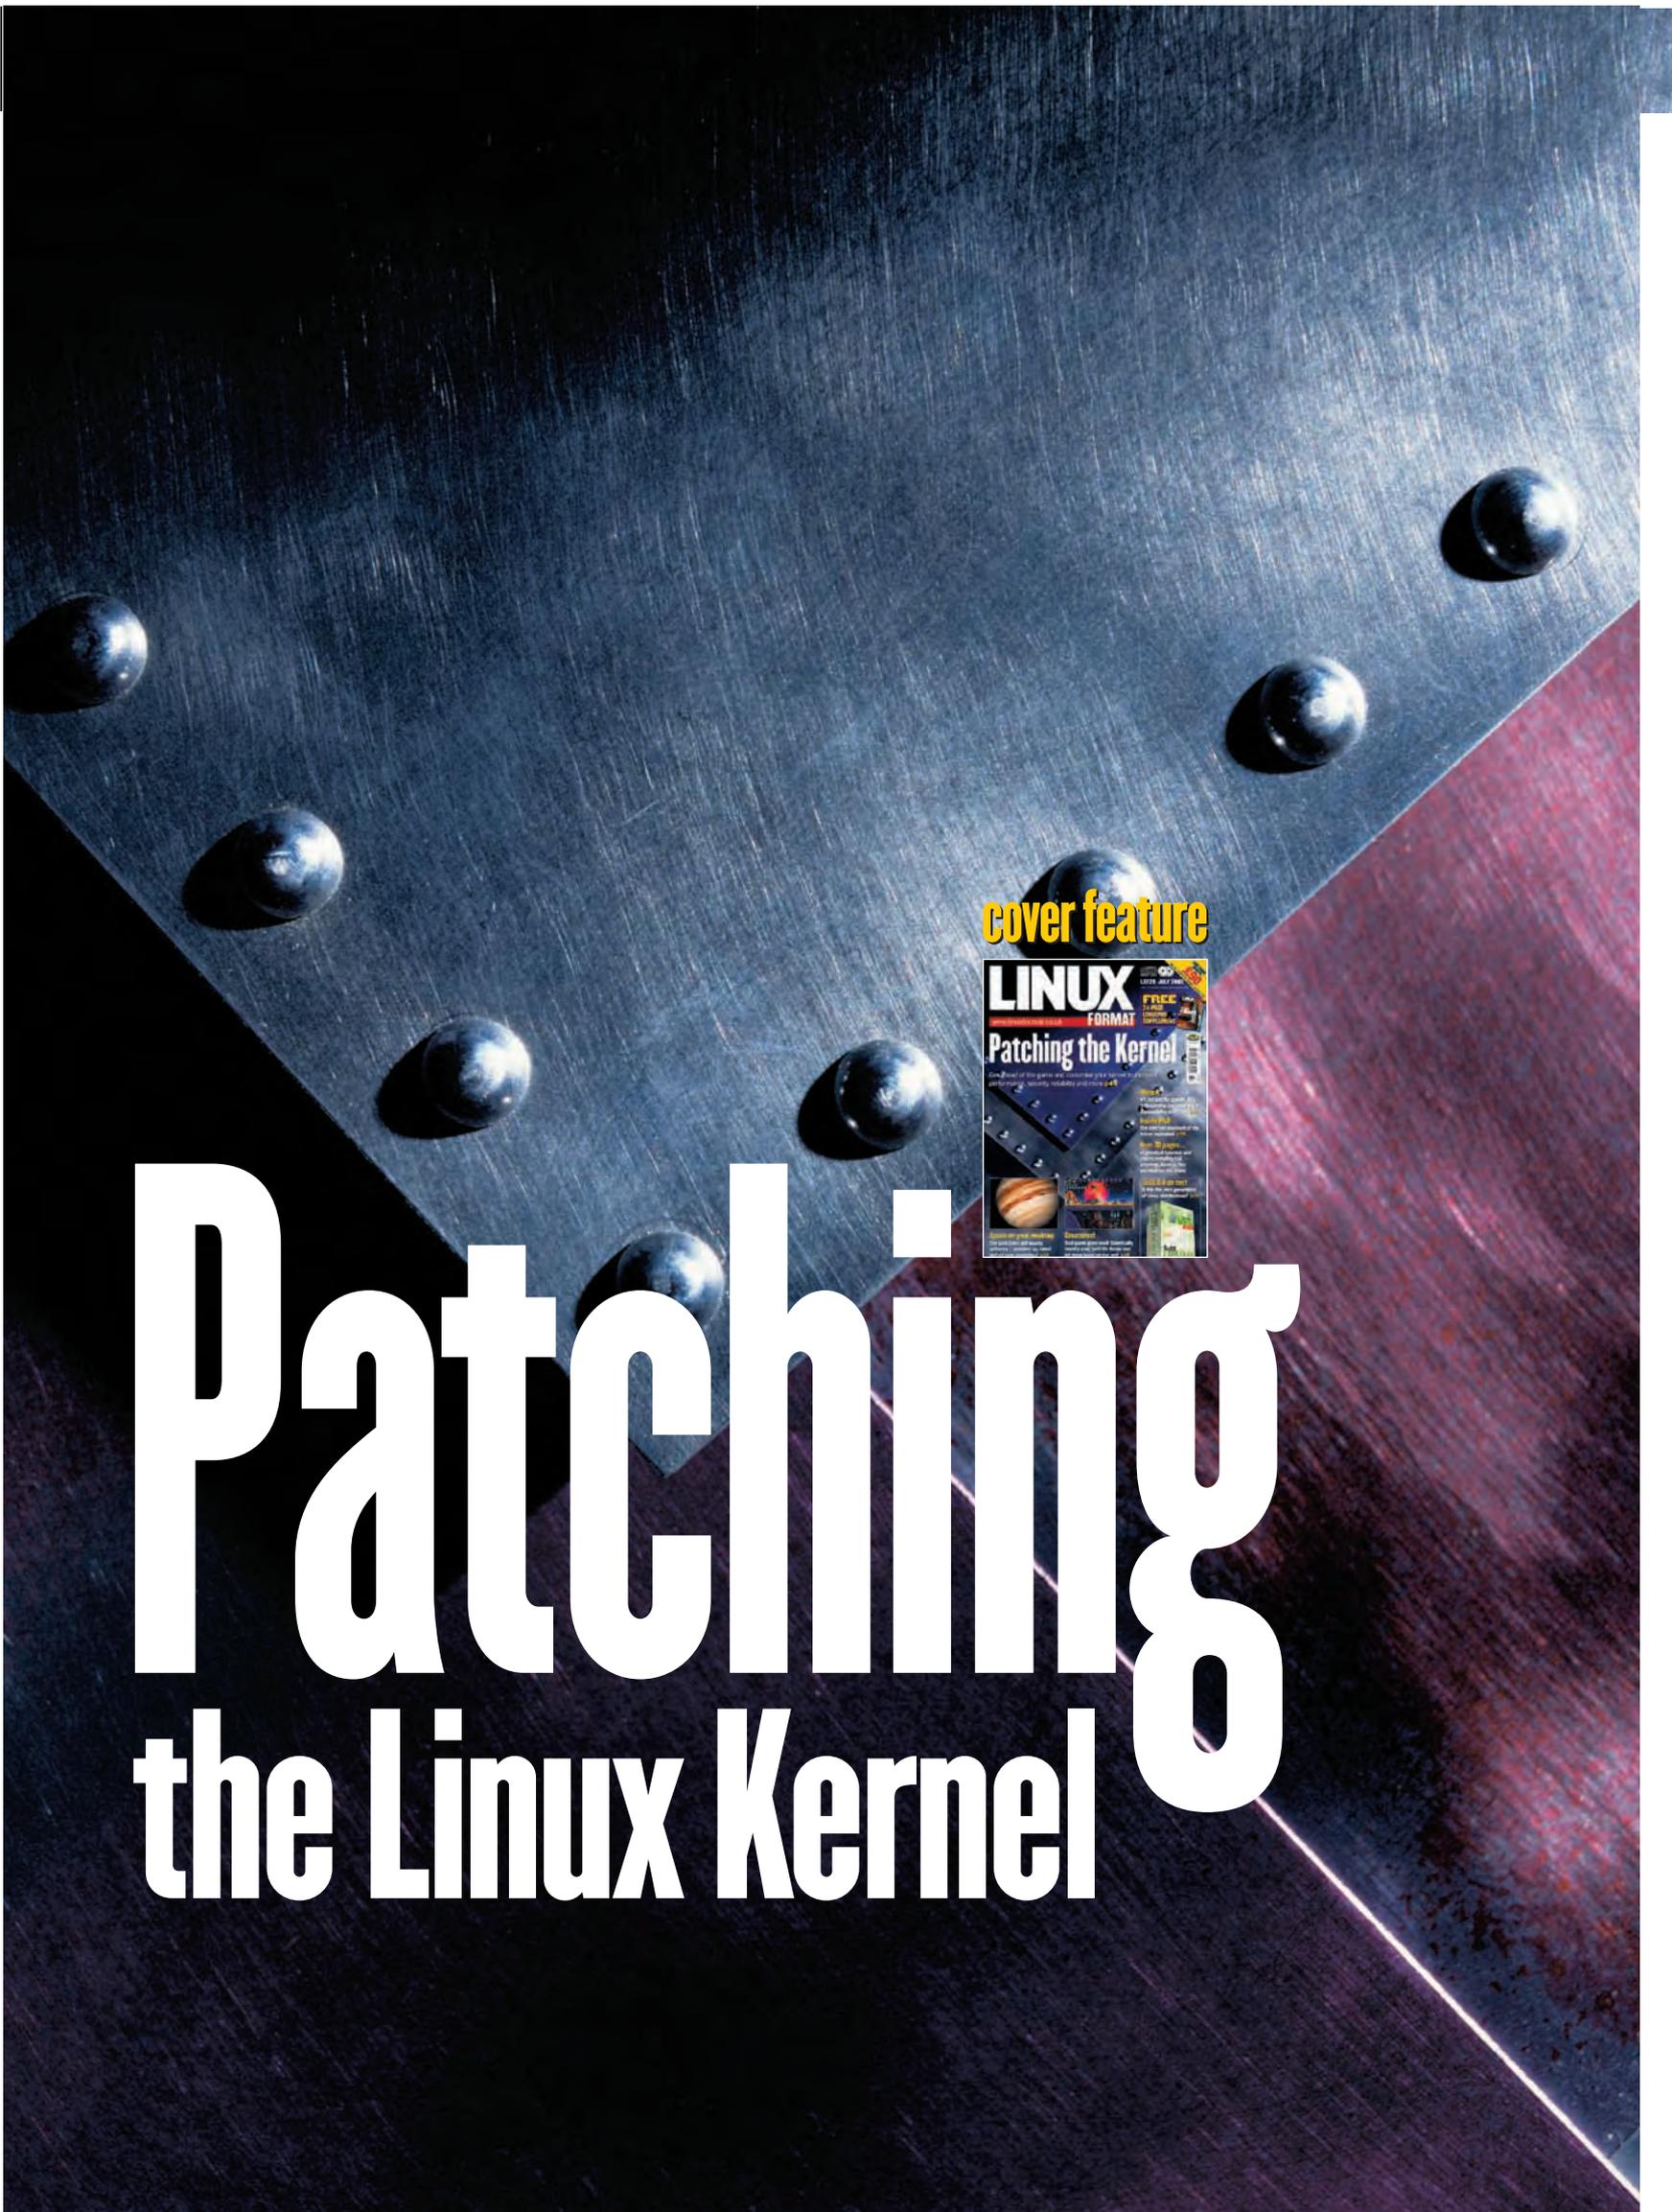
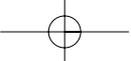


**Many X extensions – such as *GLX* and *DGA* – aren't supported yet, but it can do shaped windows.**

*XDirectFB* works well in practice. It lacks some of the niceties of the full *XFree86*, but this is partly due to the immaturity of *DirectFB* itself. *DirectFB*'s multi-application core (needed to run multiple *DirectFB* applications on the same screen simultaneously) is a work in progress, and driver support is not complete. The best-supported hardware is the Matrox range, but even then *XDirectFB* lags behind *XFree86*'s Matrox driver in performance. Running an Xmark benchmark on both proved enlightening: for **XOR** raster ops, *XDirectFB* was up to ten times slower than the conventional *XFree86*. This is what you might expect, since the *DirectFB* architecture isn't designed for this sort of thing, and in everyday use it doesn't feel appreciably slower.

One nice feature of *XDirectFB* – a consequence of its architecture – is that each X window can be assigned an opacity (or is that translucency?). If you hold down the **Ctrl** key with the pointer over a window and twiddle your mouse wheel, you can slide the opacity of the window from totally opaque to completely translucent. See-through windows don't tend to have much of a practical application, they tend to be confusing when over-used, but it's certainly a neat gimmick.

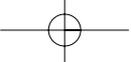
At the moment, *DirectFB* and *XDirectFB* are only really of interest to developers. *XDirectFB* is only available currently as a series of patches against *XFree86 4.2.0* – so actually building and installing it is a time-consuming and non-trivial task. However, it is encouraging to see that developers are continuing to propose alternative graphical architectures to X. [LXF](#)



cover feature



# Patching the Linux Kernel



**Not everyone agrees with Linux Tovarld's' tree, so David Coulson looks at unofficial patches for the Linux kernel. The DIY approach is easier than you think – so read our guide to available patches, then follow our comprehensive patching HOWTO guide.**

**N**early everyone who runs Linux will, at some point, have looked at the kernel – the inner workings of the GNU/Linux operating system. Whether it is to compile support for hardware devices, or to make sure you have the latest stability and security updates, keeping track of what kernel you're running and what capabilities it provides is rather important.

To confuse matters, many distributions apply patches to their kernels which are not part of the standard kernel tree. Naturally, when you come to compile your own kernel, you notice plenty of things are missing or broken, and there is then the choice of either giving up lots of fun features – which are actually useful when you have them – or sticking to an older kernel which may be insecure or unstable for the sake of some handy capabilities.

Of course, Linux distributors need to get their features from somewhere, and many, if not all, are available on the Internet as patches for the standard Linux kernel distribution. However, not all patches work with all kernel versions, and many patches require a little encouragement to apply happily with others.

Unfortunately, this does require a basic understanding of the C programming language, but failed patches are generally easily fixed, if they are to be fixed at all, and the *patch* program is smart enough to try it's best to apply the patches correctly.

Many older patches have since been included in the standard kernel tree, but as not everyone is running

2.4 kernels (many have decided to stay with 2.2 for a while) there are often backports of the 2.4 code so that 2.2 users can appreciate the benefits of having the latest and greatest kernel capabilities without sacrificing stability.

## Performance Patches

**Pre-emptive kernel patch**

[www.tech9.net/rml/linux](http://www.tech9.net/rml/linux)

**Version: 2.4.19-pre8-1**

**Included in Linux 2.5.4-pre6**

**Status: Development**

The Linux kernel has pre-emptive capabilities for user-space processes, but kernel-space processes can still hold up the entire machine. This reduces the response times of the machine – which is particularly important for a workstation, where the user is only really interested in the response time of the applications they're using. The pre-emptive patch allows kernel-space processes to be pre-empted, avoiding kernel processes blocking up the scheduler.

The pre-emptive patch is not really suitable for servers but, if you're running a workstation, the improvement is really noticeable. Many people choose to apply the lock-break patch too, which stops the kernel locking processes in particular places, thus further improving the response times of the machine.

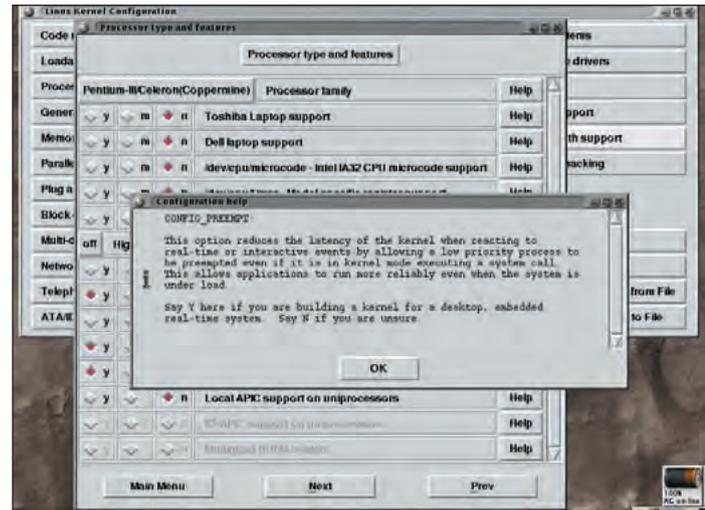
**rmap VM**

[www.surriel.com/patches](http://www.surriel.com/patches)

**Version: 13a for Linux 2.4.19-pre7**

**Status: Development**

Before 2.4 was released, there were two separate virtual memory patches floating around, both of which had their own benefits. The rmap patch by Rik Van Riel is considered by many to be a much better VM system than that included in the main Linux kernel tree, as it has improved paging and avoids large swaps which would otherwise slow down the machine. The Reverse Mapping patch is ideal for all machines where swapping is a common occurrence, such as highly-loaded servers, or low-end workstations. It is by no means perfect, but it is updated regularly and bugs are addressed particularly quickly.



**IDE Subsystem**

[www.kernel.org/pub/linux/kernel/people/hedrick](http://www.kernel.org/pub/linux/kernel/people/hedrick)

**Version: 2.4.18-rc1-02152002**

**Status: Stable**

The IDE patch is popular with Linux distributors, as it has tuning options for many common IDE chipsets for improving throughput.

Particularly if you have a separate IDE controller on a PCI card, then the IDE subsystem patches will add in much needed support for some of the more performance oriented features of such a piece of hardware.

**O(1) Scheduler**

[http://people.redhat.com/mingo/O\(1\)-scheduler/](http://people.redhat.com/mingo/O(1)-scheduler/)

**Version: K3 For Linux 2.4.18-pre8**

**Status: Development**

Using the rather popular O(1) algorithm, the kernel process

**Configuring the Linux kernel can be done either from the command line, or using a GUI tool.**

**“The popular IDE subsystem patch supports the performance oriented features of separate PCI IDE controllers.”**

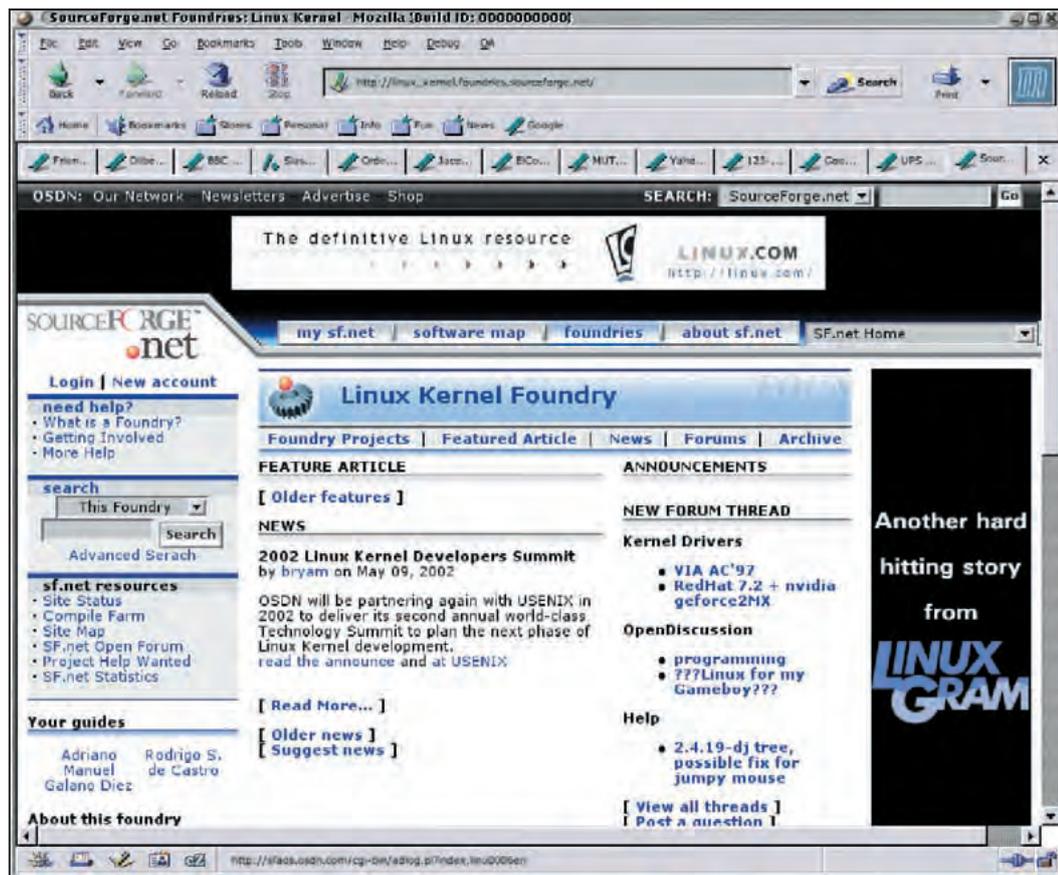
scheduler gets an overhaul, improving response times and making switching between processes smoother.

The pre-emptive kernel patch has also been modified for use with O(1) scheduler, so you can have the best of both worlds, and many people install both O(1) and pre-empt on their workstations and laptops.



# KernelPatching

Sourceforge's Kernel Foundry is a great place to start looking for Linux Kernel patches and third-party additions.



## Lock break

[www.tech9.net/rml/linux](http://www.tech9.net/rml/linux)

Version: 2.4.18-1

Status: Development

Along with the pre-emptive kernel patch, the lock breaking patch lowers latency – reducing the time spent on kernel processes by reducing the number of locks held on kernel processes.

## Productivity Patches

User Mode Linux

<http://user-mode-linux.sourceforge.net/>

Version: 2.4.18-23 for Linux 2.4.18

Status: Development

If there is one kernel patch which stands out from the rest, User Mode Linux is it. Rather than patching the kernel you run on your machine, the UML patch is applied to a separate kernel tree, producing a 'linux' binary of the kernel, which can be run from the command line from within Linux. While the UML patches are only for 2.4 kernels, you can run the resulting kernel on either 2.2 (with the appropriate patch from the UML site) or 2.4 machines without

changing their kernels at all.

UML allows the Linux kernel to run within user-space on the host, providing all the capabilities of the kernel on a physical machine, at least for anything not requiring hardware availability. Functionally, the kernel is the same, so existing binaries will run on it, and indeed, UML filesystems are distributed as distribution images from standard installations, so running Debian on UML is a realistic expectation.

Of course, UML has many uses. Probably the most common is for use as a debugging tool for Linux. Rather than rebooting the whole machine every time you make the kernel oops, you can fix the problem and then simply restart the UML instance. As one would expect, this drastically reduces the time wasted waiting for disks to *fsck* and restart everything. However, a use which is becoming more popular is as a virtual hosting platform.

UML has network connectivity capabilities, so it can access a LAN, or the Internet, through the host machine. As the UML filesystem is entirely within an image filesystem and

**User-Mode-Linux** allows Linux to execute from within Linux, creating a jailed Linux installation which can be used for development or virtual hosting.

```
david@macha:~ (pts/9)
m3Cha 18:12pm Wed May 15
david:~ ssh djc-gate, broker, davidcoulson.net
Last login: Wed May 15 21:12:18 2002 from macha, broker, davidcoulson.net on pts/8
Linux djc-gate 2.4.18-13um #23 Thu Apr 4 01:16:18 BST 2002; i686 unknown

Most of the programs included with the Debian GNU/Linux system are
freely redistributable; the exact distribution terms for each program
are described in the individual files in /usr/share/doc/*/copyright

Debian GNU/Linux comes with ABSOLUTELY NO WARRANTY, to the extent
permitted by applicable law.
Last login: Wed May 15 21:12:28 2002 from macha, broker, davidcoulson.net
djc-gate 18:12pm Wed May 15
david:~ uname -a
Linux djc-gate 2.4.18-13um #23 Thu Apr 4 01:16:18 BST 2002; i686 unknown
djc-gate 18:12pm Wed May 15
david:~ cat /proc/cpuinfo
processor       : user-mode
bogomips      : 708.97
host          : Linux echad 2.4.17 #1 Tue Feb 12 17:50:10 GMT 2002; i686
djc-gate 18:12pm Wed May 15
david:~
```

it executes in user-space, it makes for an ideal platform for running insecure services, as a test-bed for new configurations, or giving users their own 'machine' to install services, and configure as they wish.

#### IPVS

[www.linuxvirtualserver.org](http://www.linuxvirtualserver.org)

**Version: 1.0.2 for Linux 2.4**

**1.0.8 for Linux 2.2**

**Status: Production**

In many networking situations, having a single machine handle all of the load is unrealistic, so distributing requests and traffic over multiple boxes is an efficient way to reduce costs and avoid shelling out for an over-priced high-end server.

IPVS, the IP Virtual Server, is a loadable module for the *Netfilter* packet management feature in Linux 2.4, or as a kernel patch for 2.2, which controls the distribution of IP traffic entering the machine, so that it may then relay all traffic to a number of internal nodes which run the actual server, such as *Apache*.

#### Mosix

[www.mosix.org](http://www.mosix.org)

**Version: 1.5.7 for Linux 2.4.17**

**Status: Production**

Mosix is a kernel patch which allows processes to migrate to other machines on the network, so process load can be distributed over multiple CPUs. Mosix is not useful for high I/O processes, as that requires the process to be executed on the machine which started it to cut down on network bandwidth, but for high CPU usage applications, such as compiling or number crunching with SETI@home, then Mosix can do wonders if you've got a pile of machines sitting somewhere getting dusty. The only *caveat* is that Mosix is now no longer GPL, although the earlier GPL code has been salvaged and distributed as OpenMosix.

#### OpenMosix

[openmosix.sourceforge.net](http://openmosix.sourceforge.net)

**Version: 2.4.17pre7**

**Status: Development**

OpenMosix is the GPL distribution of the Mosix clustering kernel patch. OpenMosix must be installed on all systems which you want to distribute processes between.

## Patching Linux

### Our easy HOWTO

Installing patches and compiling the Linux kernel is an easy task, so here's how.

Although each and every patch is slightly different, the general rules for applying a patch to the kernel is the same whatever it may be. Patches are files containing the differences between the source tree before, and after, the changes have been made, so even once the patch has been applied, it is possible to reverse the patch and have the kernel just as it was before.

Applying a patch to the kernel tree is done using the *patch* command, and the general rule for doing so is;

```
cd /usr/src/linux
```

```
patch -p1 < ~/patch.diff
```

Usually patches are patching the file in the 'linux' directory, and often developers rename the base Linux kernel to 'linux.vanilla' or 'linux-version', so we need to use the **-p1** switch which strips off the first directory in the list. For example, if the patch is supposed to patch `linux.vanilla/fs/ext3/file.c`, using **-p1**, it will try to patch `fs/ext3/file.c`, which is why we need to *cd* into the `/usr/src/linux` directory before trying to patch it. One can also apply patches using **-p0** assuming the directory structure is as the patch expects it to be. If we neglect to pass a **-p** switch, then it will try to patch 'file.c' in the current directory, ignoring all other path information.

Patches are split into sections called hunks, and each hunk will patch a separate section of the file. *patch* will apply each hunk in order, but if one hunk fails to patch correctly, then it will continue to apply the rest of the hunks. A hunk may not patch perfectly, such as if the line number for the first line of the hunk is different from what *patch* expects, but it will try to apply the hunk anyway. If the hunk patches correctly, but needs to adjust the location of the patch in the file, then you will see this;

**Hunk #1 succeeded at 171 with fuzz 2 (offset 2 lines).**

Succeeded hunks can be ignored, as if you choose to reverse the patch, then it will do the same when it comes to remove it. However, should a hunk fail, it will return the following;

**Hunk #2 FAILED at 414.**

Once there is a failed hunk, it will apply all other hunks in the patch, then copy the original file, before any of the hunks were applied to 'filename.orig' and the failed hunks to 'filename.rej'. By looking at 'filename.rej' and 'filename' it is possible to manually apply some patches which have failed in parts. Often, failed hunks are because of struct changes in the code, or because something is formatted directly, but they can also be because large portions of the code have been changed and none of the code looks the same anymore. If you're familiar with C, then it's usually possible to apply the patch by hand anyway, assuming you know what needs to be changed in order for it to compile correctly, much less execute cleanly.

Within the patch, there are three different ways the patch can affect the file. Looking at a patch, there will be lines which start with a **+**, which means that the following text is included into the file. Patches always have text from either side of the actual code which is being changed, to ensure that the correct portion of the file is being edited, rather than having it insert lines here, there and everywhere. Some lines will start with a **-**, meaning that the line is removed from the file being patched. The final type of patch is a **!**, which means that the line should be changed to read what follows the **!**. Any individual hunk can contain any type of patch, so applying patches by hand takes a little time and thought.

Once the patches have been applied, and everything is happy, it is usually only a case of running **make oldconfig** within

the kernel tree, which rebuilds the kernel configuration using the existing `.config` file, yet steps through the configuration system and prompts for any new options which may exist, such as those provided by our patch. Once configuration has been completed, the kernel can be compiled in the normal way, and then installed. If the patch you have applied can be built as a module, it may be that you still need to recompile the entire kernel and reboot the machine, as many modules have hooks into the kernel which must be compiled in, otherwise the module will not *modprobe* correctly and instead report unresolved symbols.

Having lots of kernels on the go, some with patches, some without, it makes sense to ensure that they all have unique version names, so as to avoid confusion as to what patches are applied to the kernel you are using. Within `linux/Makefile`, there is an **EXTRAVERSION** variable, which can include extra version information. If you've applied the Pre-emptive kernel and lock-breaking patches, you might set **EXTRAVERSION** to **-pe-lb**, and then recompile it. It is often also helpful to include patch version information in the variable, as while many patches are revised once for a single kernel version, many are developed during long periods when the main kernel stays the same, so there will be more than one release of the patch for a specific kernel. Knowing which version is currently applied, and running, greatly aids in debugging and reporting problems should you experience any. Many patches are not in the kernel for a reason, and usually that is lack of stability, which will affect the entire system. As always with the kernel, a small instability can disrupt the entire system, so if you need your machine to stay up for a long time and not have problems, then you may want to wait until the patch has stabilised and has been included in the main kernel tree.

#### Software Suspend

<http://falcon.sch.bme.hu/~seasons/linux/swsusp.html>

**Version: v8 for 2.4.10**

**Status: Development**

Most laptops support suspension and hibernation capabilities, where the contents of the RAM is swapped to disk, and then resumed later.

Either on laptops which lack this capability, or on workstations which are turned off frequently, the software

suspend patch can suspend a machine and then continue where you left off last time.

The biggest *caveat* of the software suspend patch is that it currently only stores main system memory, and not them memory from any video device, so if you forget to switch to the console before suspending the system, then it will return rather corrupted which can cause X to lock-up.



# Kernel Patching

## Security Patches

### LIDS

[www.lids.org](http://www.lids.org)

**Version: 1.1.1r2 for Linux 2.4.18**

**Status: Production**

LIDS is a very popular security patch for Linux, which reduces the capabilities of the Linux kernel for specific processes. For each process, root can have different capabilities, such as the ability to edit files, bind to low ports, kill processes and so forth, reducing the opportunity for compromised services to have a major impact on the rest of the system, as the root user is not able to do any root-level operations. LIDS has been covered in *Linux Format* before, and most people choose to install LIDS on their gateway or on any Internet-facing machine which has publicly accessible services, such as a mail or shell server.

### IPSEC

[www.freeswan.org](http://www.freeswan.org)

**Version: 1.97 for Linux 2.4.18**

**Status: Production**

Security is a major consideration when providing connectivity to the

```

david@macha:/usr/src/linux-patched (pts/9)
PATCH(1)                                PATCH(1)
NAME
  patch - apply a diff file to an original

SYNOPSIS
  patch [options] [originalfile] [patchfile]
  but usually just
  patch -p<n> <patchfile

DESCRIPTION
  patch takes a patch file patchfile containing a difference
  listing produced by the diff program and applies those
  differences to one or more original files, producing
  patched versions. Normally the patched versions are put
  in place of the originals. Backups can be made; see the
  -b or --backup option. The names of the files to be
  patched are usually taken from the patch file, but if
  there's just one file to be patched it can be specified on
  the command line as originalfile.

Manual page patch(1) line 11
  
```

Internet, and when connecting to LANs over a public network, ensuring that the connection is secure, so as not to expose either LAN to external influences, is at the top of the list. IPSEC is a standard implementation for a secure IP-over-IP tunnel, which supports strong-encryption and key based authentication of the actual connection.

IPSEC is supported by many routers, including OpenBSD and Cisco, so it is not only for a Linux to Linux connection. IPSEC is rather complex, and requires careful configuration to ensure that the system is not exploitable, but once it has all been setup, it offers a very secure and capable VPN technology.

### OpenWall

[www.openwall.com/linux](http://www.openwall.com/linux)

**Version: 2.4.18v0 for Linux 2.4.18**

**2.2.20v3 for Linux 2.2.20**

**Status: Development**

OpenWall addresses many of the basic security problems with the vanilla Linux kernel, such as buffer overflows and memory access exploits. Access to /proc, /tmp and the standard file descriptors is included, as is process quota support. OpenWall is very useful on a shell server, restricting the amount of information available to users, such as the processes other users are running and system resource information which should only be available to root, rather than everyone.

### LSM

<http://lsm.immunix.org/>

**Version: lsm-full-2002\_04\_28 for kernel 2.4.18**

**Status: Production**

The Linux Security Module (LSM) is used by both LIDS and SELinux to provide security capabilities to the Linux kernel. LSM is a framework for including security modules within the kernel, rather than hacking up the kernel and stopping other things from working. LSM offers an API to loadable kernel modules. It is hoped that LSM will be included into Linux 2.5 at some stage, and become the standard method for security patches to interface with the kernel.

### SELinux

[www.nsa.gov/selinux](http://www.nsa.gov/selinux)

**Version: 2.4 for Linux 2.4.18**

**Status: Production**

While Linux is considered to be a secure platform, it has many shortcomings which can be exploited through compromised daemons, permitting remote attacks to take over a machine. SELinux has many security updates, including fine grained access control list (ACL) support, and avoids many buffer-overflow exploits.

SELinux is an ideal patch for a gateway or front-end server, and coming from the NSA, it has a well respected pedigree, giving it an edge over many other security patches for Linux.

## Filesystem patches

### ext3/jbd

[www.zip.com.au/~akpm/linux/ext3](http://www.zip.com.au/~akpm/linux/ext3)

**Version: Included from Linux**

**2.4.15**

**0.0.7a for Linux 2.2**

**Status: Production**

Journalled filesystems have been a big development point over the last few years, and ext3 has been a very popular choice, as it expands on the

The `patch` command is used to apply patches to your kernel tree.

## Branches

### The spreading Linux tree

The Linux kernel has forked many times over the last few years, so what are all of these other releases?

Performing a `finger` on `ftp.kernel.org` produces many different kernel versions. As well as the 2.0, 2.2 and 2.4 releases, there are the current 2.5 releases, along with pre-patches for the main branches. The most popular branches are those from Alan Cox, the `-ac` branch, from Dave Jones, the `-dj` branch, and Michael Cohen's `-mjc` branch which contains very bleeding-edge code for performance. The current patch from Alan Cox is `patch-2.4.19-pre8-ac3.gz`, which is the third patch by Alan, applied to the 2.4.19-pre8 kernel. However, the latest stable kernel release was 2.4.18, so how do we get 2.4.19-pre8? Within the 'testing' directory on `ftp.kernel.org`, there is a `patch-2.4.19-pre8.gz` file, which will patch a 2.4.18 kernel up to 2.4.19-pre8. Should pre9 be released, then we will need to reverse the patch for 2.4.19-pre8, then apply the pre9 patch. Going from pre8-ac3 to pre9-ac1 is a little more interesting, as you will need to remove `-ac3`, then

`-pre8`, then apply `-pre9` and `-ac1`.

Usually, the `-ac` branch is the best to track, as it contains many small patches and changes from the Linux Kernel Mailing List (LKML), which have not yet made it into the main kernel tree. Small sections of `-ac` are merged into the main kernel over time, but `-ac` is usually considered to be the most up to date kernel branch available, without sacrificing stability for strange new features which most people won't be using. Pre-release kernels are not always less stable than the main releases, and usually either small simple updates and are being worked on to produce a complete release of the kernel, so 2.4.19-pre8-ac3 should not be especially more unstable than 2.4.19.

Most patches are distributed for the latest stable kernel, the latest pre-patch if it exists, and for the current `-ac` release, so you need to be aware of what kernel you are running, and apply the appropriate patch. If you download the pre-emptive kernel patch for 2.4.18-ac1, then you will need to patch 2.4.18 to `-ac1` before trying to apply the pre-emptive patch to your kernel.

existing ext2 filesystem, which nearly every Linux user will be using, rather than starting again from scratch. In fact, ext3 is actually a patch which integrates the journaled block device with the ext2 filesystem, so the actual filesystem is still ext2 and the jbd handles the journaling side of the system. As of 2.4.15, ext3 has been included in the standard kernel distribution, so it's use is fairly widespread now.

### ReiserFS

[www.namesys.com](http://www.namesys.com)

**Version: Included from Linux 2.4.1 3.5.35 for Linux 2.2.20**

**Status: Production**

ReiserFS is one of the more mature journaled filesystems for Linux, having been in development for the last few years. It also has the rather flattering title of being the first journaled filesystem in the 2.4 kernel, although there are still a number of patches which have not been included in the kernel yet available on the site which anyone seriously using ReiserFS on a production system should look into. Unlike ext3, the ReiserFS development team have continued to develop patches for 2.2 kernels, which many consider to be more suitable for production servers which need to be stable than 2.4. Many Linux

distributions come with ReiserFS support already, and indeed, many permit you to install the system onto a ReiserFS filesystem rather than ext2

### JFS

<http://www-124.ibm.com/developerworks/opensource/jfs/>

**Version: 1.0.17**

**Included in Linux 2.5.6-pre2, 2.4.18pre9-ac4**

**Status: Production**

Coming from a highly reputable and well known company such as IBM, JFS has been given a great deal of press and is used on IBM's Linux servers, as you would expect. JFS is distributed under the GPL, avoiding the need for closed-source binaries tainting the kernel. It is, however, not quite as well used as either ext3 or ReiserFS, but it is, as you would expect, rather well maintained and IBM have a very active mailing list encouraging users to become involved in the development of JFS, so as to avoid the feeling that it is coming from a faceless corporation.

### XFS

<http://oss.sgi.com/projects/xfs/>

**Version: 1.1 for Linux 2.4.18**

**Status: Production**

XFS is the popular journaled filesystem from SGI, which has been

used in the past on their high-performance servers, and has now been ported to Linux, distributed under the GPL. Many people regard XFS as the filesystem for use on large SANs or RAID arrays, as it handles large files and large filesystems better than many other filesystems which are available.

### Online ext2 resize

<http://www-mddsp.enel.ucalgary.ca/People/adilger/online-ext2/>

**Version: 1.1.17 for 2.4**

**Status: Development**

Both XFS and ReiserFS allow the resizing of mounted filesystems, and the Online ext2 resize patch permits mounted ext2 or ext3 filesystems to be resized. Of course, to make full use of it, the system must be running LVM, or make use of a live DOS partition resizer, as the patch only handles the actual filesystem, rather than partitions or logical volumes.

### NTFS

<http://linux-ntfs.sourceforge.net/>

**Version: 2.0.7 for Linux 2.4.18**

**Included in 2.5.16**

**Status: Development**

The standard 2.4 kernel currently has version one of the NTFS filesystem, which supports reading and rather experimental writing, so version two has been a complete rewrite, offering improved speed of reading. Unfortunately, it does not yet support writing, but hopefully it will have somewhat more stable writing capabilities than previous versions of the driver.

### DevFS

[www.atnf.csiro.au/people/rgooch/linux/kernel-patches.html](http://www.atnf.csiro.au/people/rgooch/linux/kernel-patches.html)

**Version: Included in Linux 2.3.46 v99.21 for Linux 2.2**

**Status: Production**

DevFS is an alternative filesystem for the /dev directory, avoiding a messy and disorganised device directory, and avoiding having device files stored on the / filesystem of the machine. Rather than creating all of the possible /dev entries, modules register them with DevFS, which creates them on the fly, so if you don't have a piece of hardware, or don't have support for it available to the kernel, then there will

ftp.kernel.org contains kernel patches and releases from a whole selection of developers.

The release info for the kernel is stored within the Makefile.

```
david@macha:~ (pts/9)
ncftp / > cd pub/linux/kernel/people
ncftp /pub/linux/kernel/people > ls
message  aboe/      dreakh/  jes/      mingo/    rgooch/    tgall/
acne/    bcr1/      fchy/    jgarzik/  ml/       rial/      tursu/
aeb/     cort/      hadrick/ jmerkez/  mic/      rnl/      whaues/
alan/    davei/     hpa/     juhauff/  mochel/   rusty/    wll/
andrea/  daven/    hur/     laredo/   paulus/   sct/
astor/   dmh2/     jbglaw/  marcelo/  prumpf/   tao/
ncftp /pub/linux/kernel/people >
```

```
david@macha:~/usr/src/linux-patched (pts/9)
VERSION = 2
PATCHLEVEL = 4
SUBLEVEL = 18
EXTRAVERSION = -pre9-01K3-preempt-slb-ide01192002-lvm103
KERNELRELEASE=$(VERSION).$(PATCHLEVEL).$(SUBLEVEL)$(EXTRAVERSION)
ARCH := $(shell uname -m | sed -e s/i.86/i386/ -e s/sun4u/sparc64/ -e s/arm.*/*
n/ -e s/sa110/arm/)
KERNELPATH=kernel-$(shell echo $(KERNELRELEASE) | sed -e "s/-//g")
CONFIG_SHELL := $(shell if [ -x "$BASH" ]; then echo $BASH; \
else if [ -x /bin/bash ]; then echo /bin/bash; \
else echo sh; fi ; fi)
TOPDIR := $(shell /bin/pwd)
HPATH = $(TOPDIR)/include
FINDHPATH = $(HPATH)/asm $(HPATH)/linux $(HPATH)/scsi $(HPATH)/net
HOSTCC = gcc
HOSTCFLAGS = -Wall -Wstrict-prototypes -O2 -fomit-frame-pointer
"Makefile" [readonly] 542L, 17262C 1,0-1 Top
```

# Kernel Patching

The Linux Progress Patch makes Linux a little more friendlier to those who are scared by the boot up messages.



be no /dev entry for it. This is especially useful for hot-plug devices, including those using USB and IEEE1394 interfaces, and also for loadable modules, such as CDROM drives, network devices and PCMCIA media. Devfs is used by many distributions, and as it is now included in the 2.4 kernel, there are few reasons to use the regular /dev directory when devfs is available instead.

## Misc Patches

### Linux Progress Patch

<http://lpp.freelords.org/> - Themes  
<http://lpp-themes.sourceforge.net/>  
**Version: For Linux 2.4.12**

#### Status: Production

As fascinating as boot up information is, many inexperienced users are scared by it, rather than finding it to be useful, so the Linux Progress Patch was created to make the whole boot procedure a little nicer on the eyes. It is not completely useless, as it does give status when loading services and daemons, but it only shows what is really needed. With a selection of themes, to suit everyone's taste, or lack thereof, then it can make everything much prettier and simpler for less experienced users who just want to get things done.

### BadRAM

<http://rick.vanrein.org/linux/badram/>  
**Version: v1 for Linux 2.4.16**

#### Status: Production

On many older systems, the memory can become unstable or otherwise unusable because a small portion of it has been damaged or is simply worn out. If the rest of the memory board is fine, then it seems rather unfortunate that you would have to replace the entire board just to enable the machine to work properly. The



BadRAM patch forces the kernel to skip certain memory addresses, so that you can keep a less than perfect DIMM in the machine. Using the *memtest86* program, it's fairly straightforward to find out which memory locations have problems, and it will even spit out the appropriate line to pass to the kernel in order for the BadRAM patch to ignore the section of memory which is unusable.

### International Kernel Patch

[www.kernel.org](http://www.kernel.org)  
**Version: v1 for Linux 2.4.18**  
**Status: Development**

Because of export restrictions in the US, regarding encryption technologies, the Linux kernel cannot contain any encryption code, which means that for encrypted networks and filesystems, a third-party must provide the capabilities. IPSEC provides for encrypted network connectivity, but for encrypted filesystems, the International Kernel Patch enables a user to mount an encrypted imagefile over the loopback device and write to it as a normal filesystem. Using a passphrase, the user can stop anyone else accessing their filesystem, thus ensuring that whatever they store in it is secure. The use of encrypted filesystems, using the International Kernel Patch is covered at <http://encryptionhowto.sourceforge.net/>

### 802.1Q VLAN

<http://scry.wanfear.com/~greear/vlan.html>  
**Version: 1.6 for Linux 2.4.14 and above**  
**Status: Production**

802.1Q is a standard protocol for virtual-LANs, which is supported by many managed switches and routers, allowing a single LAN to be chopped up into many smaller networks.

If you don't have a monster switch from someone like Cisco or 3com, you probably won't want this patch, but using VLANs on a large switch allows machines to be grouped together, and will only broadcast to machines on their own VLAN, which avoids wasting bandwidth and machines seeing packets which have nothing at all to do with them.

### Compressed Caching

<http://linuxcompressed.sf.net/>  
**Version: 0.23pre2 for Linux 2.4.18**  
**Status: Development**

Paging physical memory to a swap partition on a disk is a slow process, which is undesirable in almost every situation. The Compressed Caching patch changes the organisation of the Linux memory hierarchy, creating two portions of memory where there used to be one.

One of these is used to store regular pages as memory, as normal, yet the other is used to store compressed pages, as opposed of swapping them to disk. This increases the size of the physical memory available to processes, which reduces the amount of overhead from swapping pages out to disk, then back into memory. There is a slight overhead from compressing pages of memory, but it is negligible compared to that of writing to a slow hard disk. [LXF](#)

## Links

### Useful websites for kernel distribution

[www.kernel.org](http://www.kernel.org) - Central kernel distribution centre. [www.uk.kernel.org](http://www.uk.kernel.org) is the UK mirror of the site.  
[www.linuxhq.com/](http://www.linuxhq.com/) - Comprehensive resource containing all kernel revisions, changelogs, and branch information  
[http://linux\\_kernel.foundries.sourceforge.net](http://linux_kernel.foundries.sourceforge.net) - Linux Kernel Foundry on SourceForge.net, with all of the SourceForge projects which are concerning the Linux kernel  
<http://linux-patches.rock-projects.com/> - A small selection of kernel patches

**WineX**

# THE TRANSGAMING revolution



70cl

100%

## Is Wine really the future of Linux games development? And what do Transgaming bring to the party? Nick Veitch has the answers.

**W**ine has been around for quite a while, and its development has seen many great leaps forward, and some not so progressive moments. It has had strong associations with Corel in the past, who leveraged the technology to launch a raft of tools (*Corel Draw*, *PhotoPaint*, *Word Perfect Office*) for Linux. Even Microsoft have taken notice of the technology, by adjusting their EULA document for Windows XP seemingly to try and legally halt its use. It has also stirred up ferocious debate amongst the Linux community. But it also has some strong allies.

*Wine* started back in 1993 as a project to help Windows 3.1 applications run on Linux. It isn't an emulator, but a layer that sits between the program code and the underlying Linux system, which attempts to fulfill the needs of software making calls to Windows APIs. As it doesn't emulate the hardware or processor, you will need a

Windows-compatible processor (*i.e.* x86) to run Windows software.

*Wine* is, in effect, a bunch of libraries which do the same for Linux as their counterparts on the Windows system. When the application makes an API call, *Wine* steps in to respond. Most of the issues with software using revolve around the user interface and system resources – *e.g.* drawing a Window, accessing a disk, making a TCP request, etc.

For the most basic of applications, this is plenty of support. So your simple text viewers aren't going to be a problem. A lot of Microsoft software proves to be more problematic because, since they wrote the Windows API, they often use shortcuts and undocumented features which are very difficult to mimic in *Wine*. Then there are problems caused by the fundamental differences between Windows and libraries – *e.g.* Windows DLLs (dynamic linked libraries) have a preferred base address, and have a special fixup table when they need to be relocated.

### Running games

Games provide several greater challenges for the *Wine* project. In the first place, unlike most other software, games are usually very intensive applications that use various methods and shortcuts to 'take over' your hardware. Most of these are endorsed by Microsoft and some features of Windows are almost specifically oriented to games programming – *DirectX* for example. Then of course,

games generally don't use the same sort of GUI interface as most windows apps, and often include their own highly customised interfaces, which rely on various different Windows resources. As these may not be the primary focus for the *Wine* project, they are supported to various levels of completeness at the moment. You can bet that any custom GUI though uses the full extent of the capabilities available to it.

Before you even get to running the game, you have to install it though, and this is where *Wine* can fall down. There are many different types of installers, which all work in different ways. If you thought the installer was just a simple script you are much mistaken – Windows installers are very complex pieces of code that can easily require as much effort and skill to get working reliably as writing a small application. It's very annoying to know that an application works, but you can't install it because the install software doesn't. Another problem is copy protection, usually linked to the way CDs are mounted on the system.

Transgaming's focus for *WineX* then is to develop the *Wine* technology in areas which have previously been overlooked or put off because they were not considered part of the core goals of *Wine*.

### Non-gaming benefits

Gaming is probably one of the major factors behind hardware development on the x86 platform. The fastest, newest processors, graphics and

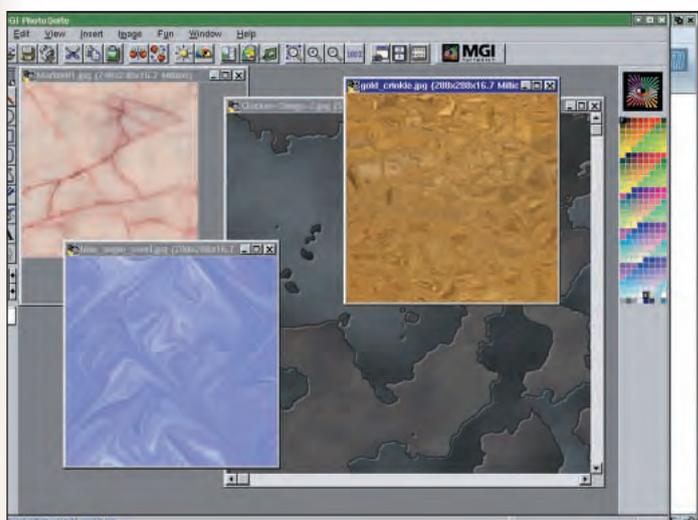
## What is Codeweavers Wine then?

### A vintage selection

The Codeweavers team have done a lot of work on the *Wine* project, and have also produced several wine based technologies including *CrossOver Office* and *CrossOver plugin*. These last two are proprietary software that allow MS Office and Windows plugins to be run on the GNU/Linux OS.

Codeweavers also produce a packaged version of *Wine*, which has a proper installer and configuration utility, and is much easier for the novice Linux user to set up and use. This version of *Wine* is straight from the main *Wine* development tree, but due to the time taken to package it up, etc, is not the very latest version of the software.

You can find out more about Codeweavers' contributions to *Wine* and their software at [www.codeweavers.com](http://www.codeweavers.com)



*DirectX* and installer technologies supported by *WineX* means more applications can run on Linux. Here's MGI *PhotoSuite*

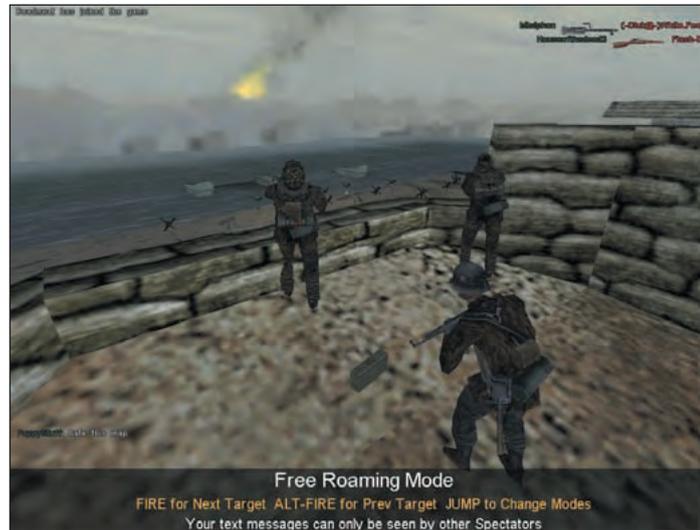


*WineX* advances mean that even games which rely heavily on *Direct2D* are playable on Linux systems.

# WineX



A number of games, like *Tropicico* shown here, will run but won't install properly on *WineX*. Yet...



*Half-life* and many of its popular modifications run on Linux, even for network play! *Day of Defeat* is well worth the download!

◀ sound cards are usually to be found in the hands of extreme gamers. Usually the games have caught up with the hardware in a few months and then everything moves on yet again. In the same way that games make extreme demands on the hardware, they also often make extreme demands on the underlying software too. So if you can get a game to work, especially a cutting edge game, you can usually get other perhaps less exciting but perhaps more useful application to work.

One of the immediate benefits of *WineX* for application users is the support for many installers. Software which may previously have worked okay, but would never install has now become available for use.

*DirectX* is also very important. Many graphics tools use *DirectX* to speed up 2D and (if appropriate) 3D rendering. Without *WineX*, they just wouldn't work.

## Just an excuse?

Many detractors of *Wine* have complained that, far from helping out Linux users, *Wine* is actually giving people an excuse not to create Linux software. Company X can't be bothered to make a Linux

version of their product if it may be possible to use the same code, saving large amounts of money and development time.

*Wine* has been responsible for some important software ports to Linux though. Corel made a big thrust into the Linux marketplace when *Corel Draw* and *WordPerfect Office* were ported to Linux, something that was possible in a matter of months rather than years thanks to the *Wine* libraries. It may have been possible to create more efficient and Linux friendly code by re-writing everything from scratch, but it just wouldn't have happened. By

making it possible to run this software under Linux, a market is created, which will encourage more native development – at least that's what the pro-Winers believe.

Whatever side of the *Wine* debate you are on, you must agree that *WineX* has made significant advances into cross-OS compatibility. Now if only I could get *Day of Defeat* to work through the company firewall, I'd be sorted...

## Licence changes

### Wine changes mean changes to WineX too

Earlier in the year, changes were made to the licence for *Wine*. Originally developed under the very permissive X11 licence, the *Wine* developers have agreed to transfer further developments of *Wine* to the GNU Lesser General Public License (LGPL).

This has had some impact on Transgaming's strategy, as well as that of others co-operating with the *Wine* project. In a briefing on the topic, Gavriel State, founder and CEO of Transgaming has said

"For TransGaming, specifically, there are a number of implications of the license change. First of all, *WineX* supports copy-protected software – almost every game available for Windows today relies on some form of copy-protection, and it is a necessary requirement for many of our partner companies. The LGPL would dictate that we publicly release the source code to our copy-protection support - an action which would violate the tenets of the US Digital Millennium Copyright Act

(DMCA), as it would divulge critical knowledge about how such technologies work. Thus, revealing this information is not an option.

"The LGPL would also lock TransGaming into the long-term support of *Wine* even in a situation where our subscription objectives are not fulfilled. This imposes non-reciprocal benefit whereby the *Wine* project is furthered by TransGaming, but our corporate objectives and the needs of our users, partners, and customers are not.

"Given these considerations, we have chosen to adopt an alternate approach, architected by TransGaming's Ove Kaven, and another long-time *Wine* developer, Eric Pouech, who are collaborating with a good cross-section of the *Wine* team to begin the "ReWind" project: <http://rewind.sourceforge.net>. ReWind is a continuation of the "X11" licensed version of the *Wine* tree and has received the endorsement of many *Wine* developers who are happy to have their changes applied to both ReWind and the LGPL'd *Wine* project."





# ON THE RECORD

We managed to catch up with Gavriel State, founder, CEO & CTO of Transgaming Technologies, and pose a few questions about WineX, Linux gaming and the future of Wine.

**LXF: Why was WineX formed as a fork from the main Wine development?**

**GAVRIEL STATE:** The *WineX* tree contains substantial enhancements that are quite specific to games, whereas development in the standard *Wine* tree is less focused. By paying attention to this narrower subset of software, we can achieve better results without worrying about unrelated modifications to the tree affecting the performance of the programs our customers care about. The *WineX* tree is also distinct from the WineHQ development tree due to differences in the licensing policy between the two trees – *WineX* is under the AFPL (Aladdin Free Public Licence), which restricts commercial redistribution of the code.

**LXF: What is the ultimate goal for WineX?**

**GS:** Simply to be the world's quintessential seamless portability solution for bringing Windows games software to new platforms. This is a utopic objective since we are working with a moving target, however, we have already demonstrated how quickly TransGaming and *WineX* can keep up: when we debuted *WineX* in October 2001, we supported only six games. Now, just four months later, we support over 80 of the most popular Windows games.

**LXF: Obviously, as well as games, the work done on WineX also helps support non-games software too (e.g. InstallShield support etc)?**

**GS:** This is correct. In order to make *WineX* a user-friendly and complete solution, a great deal of work has been done on *InstallShield*, DCOM, and other areas. This results from our desire to make *WineX* the simplest,

most seamless, and easiest-to-use portability layer available. Our objective is to serve our customers in the richest manner possible. This means that *WineX* must be transparent so that anyone, technical and non-technical customers, can use it effectively. Some of our enhancement work also supports other business software applications.

**LXF: There has been some concern in the community about forking of the WINE project. Is there any intention to fold the code back into Wine in future?**

**GS:** TransGaming already contributes large amounts of code back to the WineHQ tree. For example, we have contributed the majority of our 2D *DirectX* work, as well as substantial work on COM and OLE. We'll continue to contribute code in the future in accordance with *Wine's* own evolution.

**LXF: Given TransGaming's commitment to eventually make all of its work freely available, what will the company then do with Wine?**

**GS:** TransGaming is a strong *Wine* supporter and we are going to continue to work with the Wine community in a manner that makes the most sense for everyone. We are currently studying how best to move forward given the license changes that have taken place in *Wine* recently. (LXF Note: see box on page 54)

**LXF: WineX has been criticised by some games developers as actually being a retrograde step – instead of encouraging developers to create Linux games, it is almost giving them an excuse not to. How would you answer these claims?**

**GS:** On the contrary, as most people know, the process of porting a game is

very expensive, let alone developing games from scratch. If companies are going to invest the huge dollar amounts necessary to create a game, it makes the most business sense to pursue the largest markets. While the consumer Linux market is growing rapidly, it has not yet reached the critical threshold needed to entice games companies to do any kind of solid development work in Linux on their own. TransGaming is solving that fundamental problem; we are bringing as many popular games to Linux as fast as we can in an endeavour to increase the proliferation and utilization of Linux on the desktop. Once there is sufficient support and a demonstrable market, traditional game developers will start to add Linux to their portfolio of platforms to develop for. In the interim, Linux users want to play games and TransGaming is fulfilling that need.

**LXF: Other systems that try to co-operate with MS protocols, e.g. Samba, often find themselves hostages to development they have no control over. Isn't trying to create DirectX etc support for Linux just another case of a technology that will work for a while, but will become broken by subsequent releases?**

**GS:** TransGaming has already demonstrated that we can keep pace with *DirectX* evolution. Again, we have over 80 games running on Linux in a completely seamless manner. In December, our subscribers told us that more than anything else; they wanted us to pursue *DirectX 8* support and compatibility. In less than 10 weeks, we had our first *DirectX 8* sample programs running, and shortly after we had one of the hottest games available, *Max Payne*, fully functional.

**LXF: TransGaming subscribers get to vote on where the development effort goes, e.g. on supporting specific games, or technologies.**

**How is this system working out?**

**GS:** Very well. In fact, our voters influence our development, selecting the titles and technologies they want TransGaming to work on, and we do our best to oblige them. *DirectX 8* and *Max Payne* exemplify this commitment to our subscribers.

**LXF: The voting system allows negative votes, but doesn't this mean that individual games always get negative results?**

**GS:** We allow negative votes in order for subscribers to express the difference between indifference to a title or technology and active dislike. Our users know that every decision they make is about where our resources are directed. Some individual games receive negative votes, certainly – but that just makes it more evident which games our subscribers really want us to work on.

**LXF: The financial model for TransGaming seems to follow the "Street Performer Protocol", i.e.: the code is essentially free, but continued development depends on contributions. Is this working as expected so far?**

**GS:** This model is working out well so far and our subscribers are very committed to TransGaming's continued success since we are working very hard to serve their needs. The continued evolution of the code certainly is contingent upon the continued support of our subscribers and all Linux users in general.

**LXF: Do you have other forms of investment apart from subscribers?**

**GS:** TransGaming is working on very unique technologies and applications that solve a strong need in a variety of different markets. As a result we have received strong support from the investment community who are committed to seeing us grow.

**LXF: How many subscribers do you have currently? I gather the target was 20,000.**

**GS:** The target is 20,000 subscribers

and we are very pleased with the tremendous support we have received from the gaming community. Although we cannot disclose the exact number quite yet, we are satisfied with the growth of the subscriber base in such a short time.

**LXF: How many active developers are contributing to WineX?**

**GS:** TransGaming has assembled a truly outstanding team of developers who are also long time *Wine* contributors, in addition to the overwhelming support we receive from the *Wine* community. Our top-notch engineers come from all over the industry and possess a unique depth of Linux experience. Our team is over a dozen and growing rapidly.

**LXF: Is the feedback from subscribers/users genuinely helpful?**

**GS:** It's so helpful and important that we are incorporating a new subscriber element to our website very shortly as a gesture of appreciation to our existing members and an invitation for users to become subscribers. It will include some very entertaining and innovative contests, feedback opportunities, user profiles, prizes and more. Watch our website for up-to-the-minute information. This project is designed for more interaction with users and will feature some of their most inventive ideas and solutions. We're really excited about launching this process and to see what additional gems of wisdom and wit will come from it.

**LXF: To date there is only one 'officially' supported title, *The Sims*. Will there ever be any more?**

**GS:** First, let me explain 'official' support. *The Sims* for Linux is a title that TransGaming has exclusive rights to and our *WineX* technology is integrated into the original source code enabling it to run on Linux. *The Sims* for Linux has been modified in several ways to make it more compatible with Linux users' needs: the sound system uses the free-software Ogg Vorbis file format for playing music, the HTML export feature is compatible with Linux browsers like *Konqueror*, and video performance was carefully tuned for

running with *WineX*.

We also have over 80 games, made exclusively for Windows, which run on Linux using our downloadable *WineX*, available to our subscribers. Many of these titles are still officially supported by TransGaming, in spite of not having licensing agreements with the original developers, and include such games as *Starcraft*, *Baldurs Gate 2*, *Sacrifice*, *Half-Life* and *Counter Strike*, among others.

These approaches, although overlapped in certain ways, are distinct. There will be other titles that we will be selling soon, where TransGaming is working directly with the original developers. We utilize both models because we want to gain a greater degree of presence and

**"If companies are going to invest huge dollar amounts to create games, it makes sense to pursue the largest markets."**

strength within the game development community while ensuring our users have an abundance of games to play.

**LXF: Has there been any interest in supporting WineX from the commercial game developers?**

**GS:** TransGaming is generating a lot of interest in the gaming world and we are currently in discussions with a number of companies regarding increased support of *WineX*. One of the primary reasons why TransGaming was able to obtain an exclusive license to *The Sims* from EA was their interest in the strength of what *WineX* delivers.

**LXF: Some parts of WineX are licensed using the Aladdin Free Public Licence. Why choose this one?**

**GS:** The AFPL strikes a great balance between our need to derive revenue from our products and the needs of users and developers who want to get their hands dirty with source code. The AFPL restricts commercial redistribution of our code, but still allows the curious and more technically minded users to experiment with the code – quite a number of them have already been helpful to the *WineX* effort. [LXF](#)



# What on Earth is... IPV6?

Steve Heaven explains all about the biggest change to the Internet since it's inception.

## »» What is IPV6?

Internet Protocol Version 6. The Internet Protocol is the IP part of TCP/IP. It is the underlying protocol upon which TCP/IP networks generally, and the Internet in particular, are built. The current Internet Protocol is version 4.

## »» Why do we need it?

The main push for a new IP technology came from the need for more addresses. There are over 8 million websites, 1 billion users and over ten thousand ISPs worldwide. The rapid growth in users with 'always on' connections such as DSL has meant that IP addresses aren't shared between users as much as they were with dialup connections. The use of NAT techniques has given some breathing space before the crunch comes when IPv4 addresses are exhausted. With 'net connection spreading to handheld devices, set top boxes and domestic equipment such as central heating systems, fridges, cookers and even toasters, the crunch isn't far away.

As well as the address shortage problem IPv6 includes a number of improvements over IPv4 in terms of features and functionality. This is to address the perceived development of the Internet and telephony communications in the future.

## »» Is it compatible with IPv4?

Yes, The two protocols can exist side by side

on the same network. Also IPv6 has a special compatibility addressing mode to encompass existing IPv4 addresses.

## »» What are the differences between IPV6 and IPV4?

The main obvious difference is in addressing. IPv4 uses 32-bit addresses. This allows a theoretical number of  $2^{32}$  or ~4 billion hosts to be connected on the Internet. Due to 'holes' within allocated ranges and the legacy of the old system of class A,B,C and D addresses, not all values are available for use today. IPv6 addresses use 128-bit technology, which results in  $2^{128}$  theoretically addressable hosts. This is a very large number. Enough in fact for several thousand addresses for every person in the world.

## »» What do IPV6 addresses look like?

With IPv4 we have the familiar 'dotted quad' notation for representing a 32-bit IP address such as 212.113.202.71. Addresses in IPv6 are 128 bits long. Four times longer that with IPv4. These could be expressed as 16 decimal numbers separated by dots, but this would be very awkward to use. Instead they are usually divided into groups of 16 bits written as four hex digits, and the groups are separated by colons. An example is:

`fe80::02a0:4cff:f7a5:1af5`

A number of consecutive zeros can be abbreviated by a single "::" once in an IPv6 address. The above address is thus equivalent to

`fe80:0:00:000:2a0:4cff:f7a5:1af5`

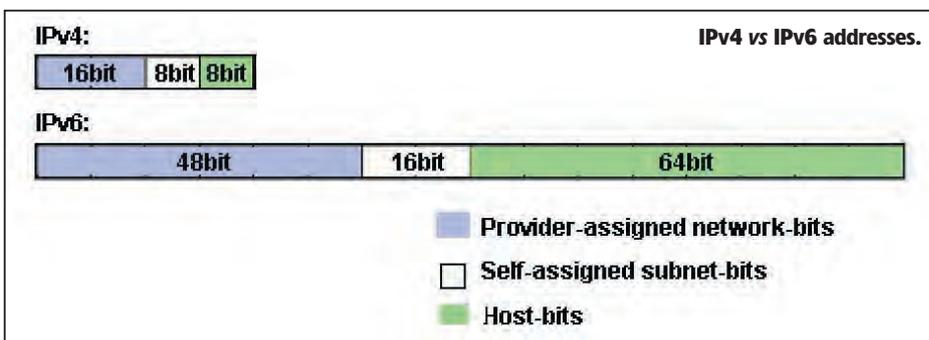
The leading zeros within groups have been omitted. As with IPv4, addresses are split in two parts – the bits identifying the network a machine is on, and the bits that identify a machine on a network or subnetwork.

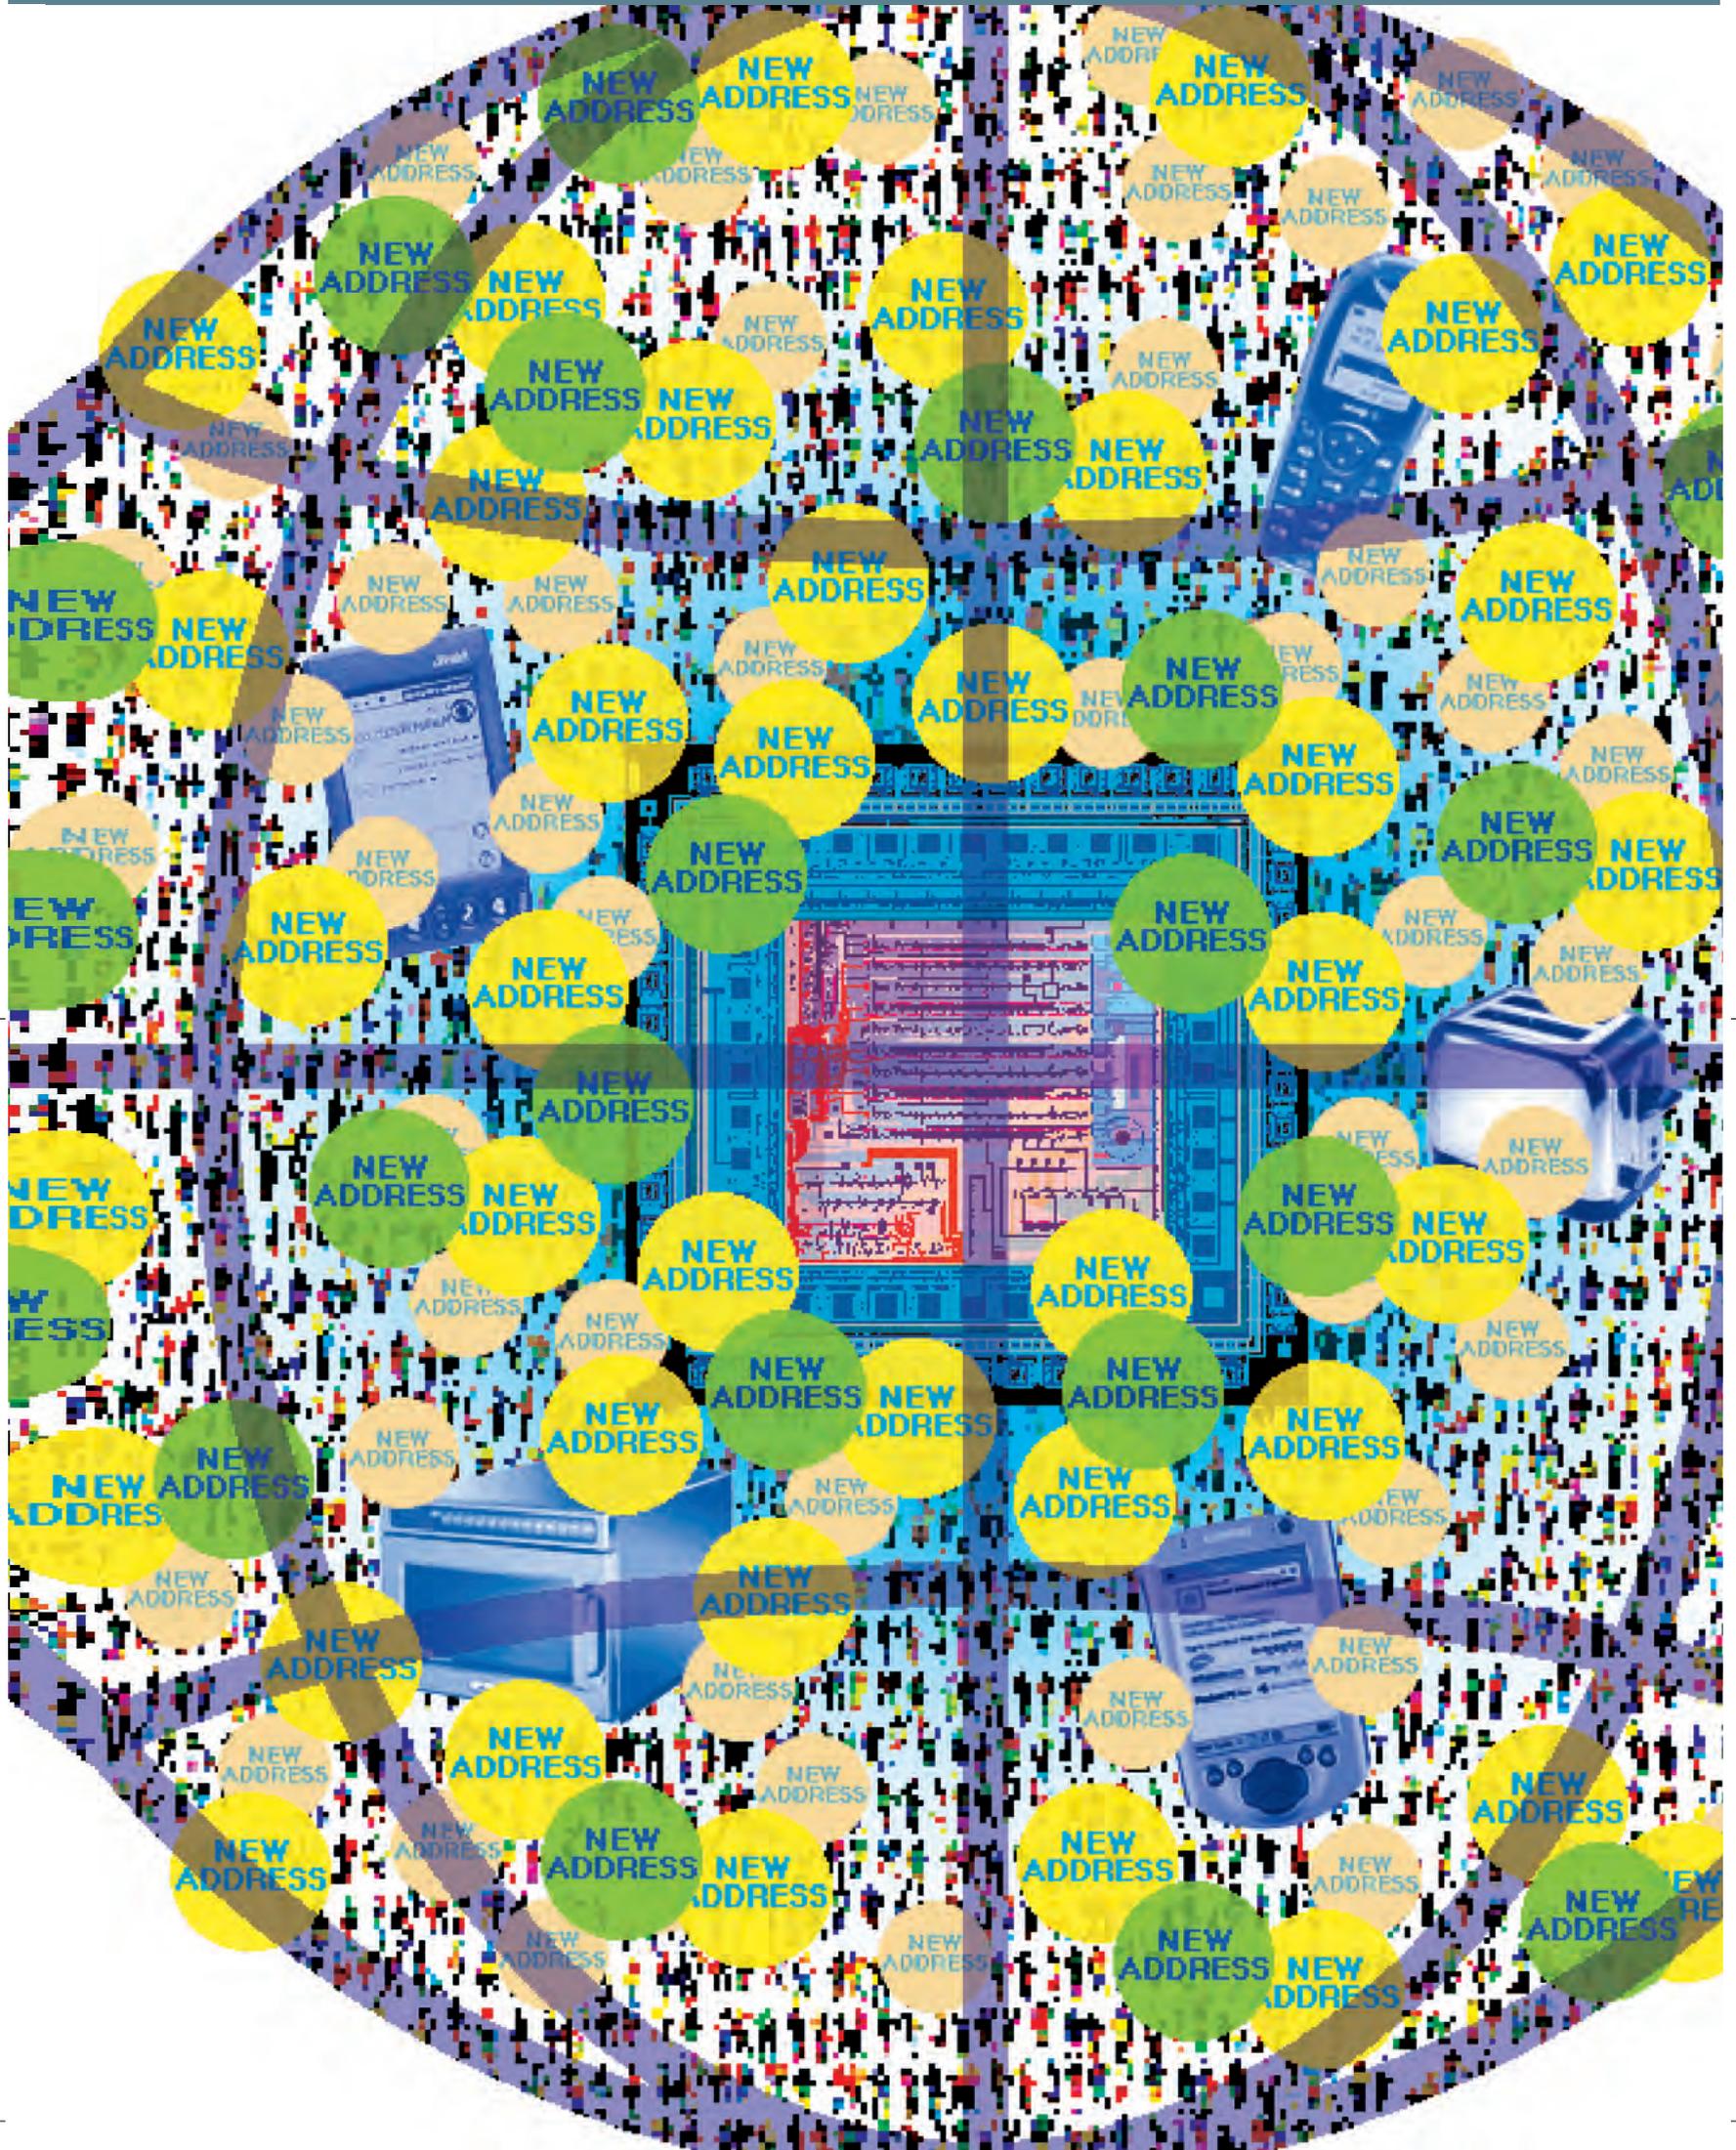
In IPv4, the border is drawn with the aid of the netmask, which can be used to mask all net/host bits. IPv4 netmasks were originally based on address classes A,B,C etc. Later with CIDR the netmask was expressed like 192.168.1.0/24 meaning the most significant 24 bits for the network part.

A similar scheme is used in IPv6: 2001:638:a01:2::/64 tells us that the address used here has the first (left-most) 64 bits used as the network address, and the last (right-most) 64 bits are used to identify the machine on the network. The network bits are commonly referred to as the (network) "prefix", and the prefix here would be 64 bits. With 128 bits available for addressing in IPv6, the scheme commonly used is the same, only the fields are wider. Providers usually assign /48 networks, which leaves 16 bits for a subnetting and 64 host bits.

## »» What are the benefits of IPV6?

IP Mobility and roaming. For use with mobile devices and IP a special protocol is defined as part of IPv6. Implementing this protocol is one of the requirements for every IPv6 stack. Thus, if you are using IPv6, then you have support for roaming





# WhatOnEarthIPv6

◀ between different networks, with global notification when you leave one network and enter the other one.

## Autoconfiguration

A scheme for assigning the host part of IPv6 addresses has been devised which will allow for autoconfiguration of addresses. The idea is to allocate fixed-width, 64-bit wide host identifiers. These need not be assigned manually as in IPv4. Instead, IPv6 host addresses are recommended to be built from so-called EUI64 addresses. EUI64 addresses are derived from MAC addresses of the underlying network interface. For example, with Ethernet, the 6-byte (48-bit) MAC address is usually filled with the hex bits "ffe" in the middle – the MAC address:

**01:23:45:67:89:ab**

results in the EUI64 address:

**01:23:45:ff:fe:67:89:ab**

which again gives the host bits for the IPv6 address:

**::0123:45ff:fe67:89ab**

These host bits can now be used to automatically assign IPv6 addresses to hosts, which supports autoconfiguration of IPv6 hosts – all that's needed to get a complete v6 IP number is the first (net/subnet) bits. IPv6 also offers a solution to assign them automatically. On a network of machines speaking IP, there's usually one router which acts as the gateway to the outside world. With IPv6 this router will send "router advertisement" information which clients are expected to either receive during operation or request on startup. The router advertisement information includes data on the router's address, and which address prefix it routes. With this information and the host-generated EUI64 address, a v6-host can calculate its IP number, and there is no need for manual address assignment.

## Multiple addresses

With IPv4, a host usually has one IP number per network interface. With IPv6, the situation is different. Each interface has its own globally unique IP address, but also has two other addresses that are of interest: The link-local address, and the site-local address. The link-local address has a prefix of fe80::/64, and the host bits are built from the interface's EUI64 address. The link-local address is used for contacting hosts and routers on the same network only, the addresses are not visible or reachable from different subnets. There is a choice of reaching the host using either the global address or using site-local address. Site-local addresses are assigned the network address fec0::/10, and subnets and hosts can be addressed just as for globally addressed networks. The only difference is that the addresses will not be visible to outside machines, as these are on a different network, and their site-local addresses are in a different physical net. As with the 192.168.0.0/16 range (see RFC1597) in IPv4, site-local addresses can be used, but don't have to be. For IPv6, it will be usual to have hosts assigned a local link and a global IP

address. Site-local addresses are not likely to be used in most situations.

## Security

Support for security was built in to the design of IPv6. IPv6 implementations are required to include *IPsec*. *IPsec* allows authentication, encryption, and compression of IP traffic. Except for application-level protocols like SSL or SSH, all IP traffic between two nodes can be handled without adjusting any applications. This means that all applications on a machine can benefit from encryption and authentication, and that policies can be set on a per-host (or even per-network) basis, not per application/service.

## Quality of Service (QoS)

IPv6 has inbuilt features to support controllable quality of service. Routers in IPv6 networks will use these to give preferential priority to some packets over others. These packets are identified by their 'traffic class' and 'flow label'. Traffic classes identify the quality of service required, e.g. a minimum bandwidth in terms of bits per second. The flow label is used together with the source address to identify a 'flow'. All packets of the flow are then given the requested QoS. These features will be used to support the delivery of video, telephony, etc over the Internet.

## ▶▶ What about Multicasting?

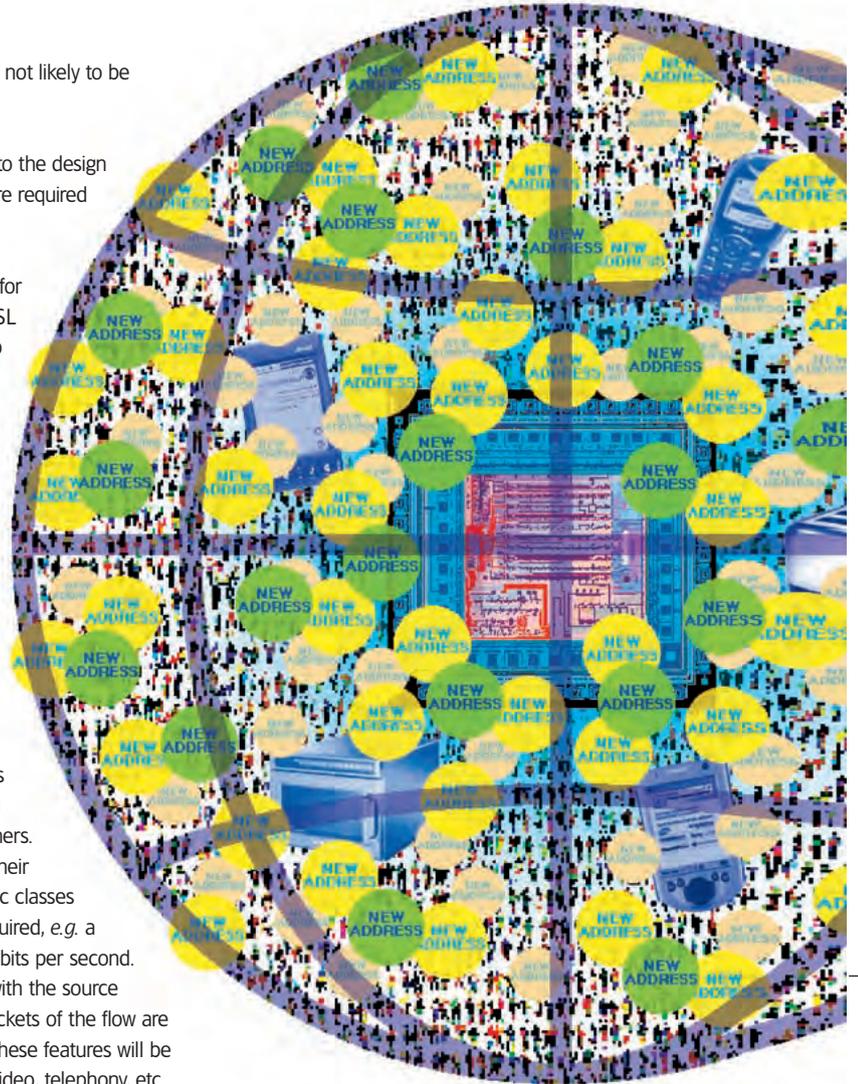
IPv6 supports multicasting much better than IPv4. It was seen as important in the design phase of the new protocol. The use of the Internet to deliver targeted broadcasting of content such as video and audio is a potential growth area. With an IP network there are three ways to talk to a host: unicast, broadcast, and multicast. The usual way to communicate with a host is by directly using its unicast address. In IPv4, the unicast address is the "normal" IP address assigned to a single host, with all address bits assigned. The broadcast address used to address all hosts in the same IP subnet has the network bits set to the network address, and the host part set to all 1s. Multicast addresses are used to reach a number of hosts in the same multicast group, which can be machines spread across the Internet. Machines must join multicast groups explicitly to participate, and there are special IPv4 numbers used for multicast addresses, allocated from the 224.0.0.0/8 subnet (see RFC1054). Multicast isn't used very much in IPv4, and only few applications like the Mbone audio and video broadcast utilities use it. In IPv6, unicast addresses are used the same as in IPv4, all the network and

host bits are assigned to identify the target network and machine. Broadcasts are no longer available in IPv6 in the way they were in IPv4, this is where multicasting is used. Addresses in the ff::/8 network are reserved for multicast applications, and there are two special multicast addresses that supersede the broadcast addresses from IPv4. One is the "all routers" multicast address, the others is for "all hosts". The addresses are specific to the subnet, for example, a router connected to two different subnets can address all hosts/routers on any of the subnets it's connected to. Addresses here are:

**ff0X::1 for all hosts** and

**ff0X::2 for all routers,**

where "X" is the scope ID of the link here, identifying the network. Usually this starts from "1" for the "node local" scope, "2" for the first link, etc. Note that it's perfectly OK for two network interfaces to be attached to one link, thus resulting in double bandwidth. One use of the "all hosts" multicast is in NDP, where any machine that wants to communicate with another machine sends out a request to the "all hosts" group, and the machine in question is expected to respond. NDP (Neighbor Discovery Protocol, see RFC2461) is the IPv6 replacement for IPv4's ARP protocol. It is used in





the autoconfiguration described earlier, but is also more flexible than ARP doing more than just IP to MAC address translation.

### »» Why hasn't the Internet changed over already?

This is a bit of a chicken and egg question. In order to change over to IPv6 a huge investment in equipment, routers etc, will have to be made by the world's ISPs and backbone carriers. Also the equipment manufacturers Cisco, Nortel, *et al.* have to invest in the development and production of the new equipment. Neither were very willing to do this until they could see a market demand. On the other hand there wasn't much demand as there are very few applications that can take advantage of or need IPv6. The situation is improving, there are products now available, but as yet they have not been deployed to a great extent except in testing and prototype setups.

### »» How will the changeover from IPv4 to IPv6 take place?

Clearly this has to be done gradually. Everyone can't change their equipment overnight. Successful trials have taken place in IPv6 Tunnelling. This is where IPv6 packets have been sent between two IPv6 hosts over an IPv4 network (*i.e.* the Internet). The technique uses special dual-stack IPv4/IPv6 routers. These routers take the incoming IPv6 packets and encapsulate them inside IPv4 packets. These then traverse the Internet as usual and are unpacked at the far end. This technique will be used in the early stages of IPv6 deployment. The opposite will be used, *i.e.* encapsulating IPv4 packets inside IPv6 packets, when most of the network is IPv6, but some older IPv4 applications remain. There is likely to be quite a long period lasting several years when both protocols are in use.

### »» What happened to IPV5?

IP V5 was reserved for ST-2 – Stream protocol version 2 for ATM, designed to control the data flow. It was abandoned in favour of RSVP.

### »» Does DNS change?

Yes. IPv6 addresses are even more difficult to remember than IPv4 ones. So of course most often a host will be referred to using a name. Simple name to address lookup maps like `/etc/hosts` are supported in the Linux resolver. The new file for IPv6

is called `/etc/ipnodes` For DNS, there are no huge changes. IPv6 name resolution is done with AAAA records that point to an entity that's four times the size of an A record. The AAAA record takes a hostname on the left side and on the right side, there's an IPv6 address, such as

```
myhost IN AAAA
3ffe:430:600:7D:ABCD:95ff:fe40:4E51
```

For reverse resolution IPv4 uses the historical `in-addrarpa` zone, subdomains below that are named using the bytes (in decimal) in reversed order. For IPv6 this is similar, only hex digits representing 4 bits are used instead of decimal numbers and the top level domain is `ip6.int`. *BIND 9.x* already understands AAAA records, so all you need to do is setup the zone files correctly.

### »» How can I run IPV6 on my Linux box?

You need a kernel version of at least 2.1.44. IPv6 support is supported in the kernel as a compile-time option. A 2.2 or 2.4 kernel has more integrated support for IPv6. To check if your kernel supports IPv6 check the `/proc/sys/net` directory. If there is an IPv6 sub directory then you have support. You will also need the latest *nettools* (*route*, *ifconfig*, *arp*, etc) package. Most modern distros include this out of the box. To check if they were compiled with IPv6 support run `ifconfig -V`. If you see a plus sign in front of `inet6` then it does support IPv6 addresses.

### »» Which distros come with IPv6 support already?

All the major Linux distros have some support in their latest versions. The minimum versions to provide some support are:  
Red hat 6.2, but you really need 7.2 to do anything  
Mandrake 8.0  
SuSE 7.1  
Debian 2.2  
Slackware 7.1

### »» What do I need to write IPv6 enabled applications?

To write and compile your own programs with IPv6 support you need either *libc5* with *libinet6*, a patched *glibc 2.0*, or *glibc 2.1*. You get *libinet6* in the *inet6-apps* package.

### »» Can I use it now?

Yes, but you need to be on an IPv6 enabled network. You can either set one up yourself from scratch or find an already connected network. Several ISPs (e.g. <http://IPv6tb.he.net/>) are offering connections to the '6bone' ([www.6bone.net](http://www.6bone.net)). This is an IPv6 testbed. The 6bone is currently a world wide informal collaborative project, informally operated by a working group of the IETF. The 6bone started as a virtual network (using IPv6 over IPv4 tunnelling/encapsulation) operating over the IPv4-based Internet to support IPv6 transport, and is

slowly migrating to native links for IPv6 transport. To help you get connected there is a quick and dirty HOWTO here: [www.linuxhq.com/IPv6/linux-IPv6.faq-5.html#ss5.4](http://www.linuxhq.com/IPv6/linux-IPv6.faq-5.html#ss5.4) Some UK universities are taking part in the trial and you may be able to persuade a netadmin to give you access. Demon and Easynet are two of the UK ISPs on the 6bone, but they aren't offering access to the public yet. You may find it useful to get the IPv6 utilities (*ping6*, *tracpath6*, *traceroute6*...). These are useful to test your setup

### »» Which applications are available?

IPv6 enabled versions of *Sendmail* and *BIND* were both released in March 2000. Many of the other major Open Source applications have released or are developing IPv6 enabled versions. All the usual suite of net apps *e.g.* *inetd*, *whois*, *finger*, *telnet*, *ftp*, etc are available, both server and client. The *Apache* web server from version 2.0.35 is IPv6 enabled. Current versions of *Lynx*, *Mozilla* and *Netscape* are OK, but not yet *Opera*. Both *MySQL* and *PostgreSQL* are working towards support and have beta versions for testing purposes. There are some patches available for *Samba*, but nothing yet in the official distribution. Patches are also available for *Postfix* and *Qmail*. The KDE desktop suite of apps has support for IPv6 from v2.1, but this is more complete in v2.2. The situation for GNOME seems more complicated. Some of the GNOME apps seem to have support, but it depends on which version of *gnomelibs* you have installed. It certainly works OK on FreeBSD, but I've yet to find someone with a working Linux setup.

### »» Can I play Quake?

Not easily :( Out of the box IPv6 enabled *Quake* servers and clients are only available for FreeBSD, Solaris and Win2000/NT currently. But if you don't mind building your own by applying various patches and fixes then sources, etc are available at: [www.quakeforge.net/search.php](http://www.quakeforge.net/search.php)  
[www.viagenie.qc.ca/en/IPv6/quake/IPv6-quake.shtml](http://www.viagenie.qc.ca/en/IPv6/quake/IPv6-quake.shtml)

### »» Where can I find out more?

The main IPv6 HOWTO is here: [www.tldp.org/HOWTO/Linux+IPv6-HOWTO](http://www.tldp.org/HOWTO/Linux+IPv6-HOWTO)  
General information on IPv6 can be found at [www.ipv6.org](http://www.ipv6.org) Some interesting presentations on IPv6 are here: [www.ipv6.ae/slides.html](http://www.ipv6.ae/slides.html) The UK IPv6 Resource Centre: [www.cs-ipv6.lancs.ac.uk](http://www.cs-ipv6.lancs.ac.uk) Various Linux IPv6 stuff can be found at [www.bieringer.de/linux/IPv6](http://www.bieringer.de/linux/IPv6)

You can join the Linux netdev mailing list at [majordomo@nuclecu.unam.mx](mailto:majordomo@nuclecu.unam.mx) or, for a general mailing list about Linux networking, subscribe to [linux-net@vger.rutgers.edu](mailto:linux-net@vger.rutgers.edu). Note that these are deeply techie groups. Only subscribe if you really want to know what's going on in the Linux IPv6 world. [LXF](http://www.linuxformat.co.uk)

# Emulators



# 6502 special

Simon Goodwin unearths seven more emulators for Commodore eight bit micros.



Your coverdisc contains *VICE*, *XC16* and *YAPE* sources and code, *Minus4* and *PfauZeh* archives, exclusive updates of *YAPE* and *XC16*, *cbmconvert*, classic Jeff Minter software, and other Commodore-related goodies.

Commodore entered 1985 with five micros, the portable *SZX-64*, (left) the seriously-styled *8296*, peak of the *Pets* (centre), the *Plus 4* (front), *C16* (top right) and classic *C64*, bottom right (see *LXF15*).

This month we have a tale of nice and nasty keyboards, custom chips, rude names in translation, and Commodore's determination to wring every drop out of the 6502. The C64, NES, Apple II and BBC Micro were famous users of the 6502 processor, as noted in past columns. Mos Technologies 6502s were at the core of Commodore's VIC-20, Plus 4 and C16, and the first one-box home computer, the Commodore PET, introduced in 1977.

## Heavy PETting

Commodore rose to prominence in the cut-throat calculator market of the 1970s, and despite a solid metal case and futuristic styling their PET had calculator-like aspects, like a built-in screen and a flat keyboard on a rectangular rather than staggered grid. The former made it the first all-in one computer, with a far better display than the TV-based micro rivals, but the chicklet keyboard was a mistake that confounded typists. It was soon replaced with a typewriter one, but not before Sharp's MZ-80K copied it.

The PET displays a monochrome grid of 24 lines of 40 characters, preset from an eight by eight grid. Many special graphic characters were included in the ROM, and represent cursor-moves and similar display functions in listings. The inbuilt BASIC was a slightly-patched 8K interpreter, bought outright

from Microsoft before Balmer and Gates caught onto the idea of selling the same thing over and over. The built-in cassette drive loaded programs at 300 baud, but the 5.25" floppy disk drive, connected to the IEEE-488 parallel user port, was capable of shifting several kB per second, with luck and a following wind.

PET stood for **P**ersonal **E**lectronic **T**ransactor, upsetting Philips, who claimed it as a trademark, and Francophones were surprised to encounter a computer called a 'fart'. Original PETs came with 4K or 8K RAM, with spare sockets drilled out to discourage DIY upgrades. The 6502 processor was designed, like the PET itself, by former Motorola man Chuck Peddle, who peddled his chip factory to Commodore, giving them a head-start in custom chips which stood them in good stead for 15 years.

One emulator for Linux mimics many PET models, with variations accommodated by command line switches and menu options. *XPET* is part of the *VICE* package, and emulates the original PET 2001, and the later 3000, 4000 and 8000 models, with 80 column displays and up to 96K RAM.

Commodore's last 8-bit business computers were the 600 and 700 series, emulated by *CBM2*. These had larger, separate screens, 128K or 256K of bank-switched RAM, expandible to 1024K as their 2MHz 6509 processors combined an 8-bit data bus with 20-bit addressing. The unreleased 630 and 730, meant to have Z80 or 8088 co-processors, are not emulated.



## VIC(s)tory

Commodore's VIC-20 was introduced in 1980, and sold millions worldwide over the following decade thanks to a good keyboard, passable sound and fast though chunky colour graphics. With just 23 lines of 22 characters on the TV screen, the small number of character cells in a complete display meant coders could obtain fast pixel-addressable colour graphics by redefining characters, when rival systems were stuck with preset fonts or mono bitmaps.

The name VIC-20 was not auspicious. The 20 refers vaguely to internal memory, but the majority of that is fixed ROM code,

derived from the earlier PET – which lacked specific support for the sound and graphics. The letters stand for **V**ideo **I**nterface **C**hip – the C64 equivalent is labelled VIC-2 – but sounds rude in countries which pronounce the letter V as an F, so the machine was known as a VC20 (Volks Computer) in continental Europe.

Only 3K of the 5K RAM is available for the user, but users soon added more – first 3K to fill a gap at address 1024 in the memory map, giving 8K vital for most commercial releases, then expansion in 8K steps up to address 32768 where the internal ROMs lurked. Further RAM or ROM cartridge expansion can stuff another 8K hole at around 40K.

## Plus 4 and C16

The successful C64 reused the VIC-20 case with revamped silicon. We tested ten Linux-friendly emulations of that system a year ago (*LXF15*). Commodore followed that with low-cost variations, mass-produced in millions in the late eighties.

The Commodore 16, Plus 4 and related machines were easier to program but failed to repeat the success of the C64. By the time they came out their minimalist 6502 processors were no match for the newer 16-bit systems, though they attracted healthy software support, and now there's a choice of Linux emulators for both C16 and Plus 4.

The C16 was an economy version with 16K RAM, but graphics boosted from 16 to 121 colours. The TED chip that did this also lacked sprites. ROM cartridges and special versions of simpler C64 games were produced for the C16, and for a while it was well supported with budget cassettes, but the addition of extra ROM holding an editor/assembler and much-improved BASIC was not enough to bring it out from under the C64's shadow.

## Transfers

### Data resurrection

You can transfer programs via the real VIC's user port and a cable to your parallel port, although you need both machines in the same room to do this, and must type in a short program at each end to manage the transfer.

Another neat way to move programs is via the 1541 disk format, given some VIC software already on old CBM diskettes. The venerable C64 drive was developed for the VIC-20, and can be connected to Linux and accessed with PD handlers written for C64 emulators. *cbmconvert* converts programs in the raw '.prg' format to the '.p00' format *XVIC* needs, which has a 24 byte header:

```
cbmconvert -n -P vic/traxx.prg
```

This creates *trax.p00* in the current directory. You can then load this into *XVIC* by supplying the path name after the emulated drive number, e.g. *xvic -fs8 ~/cbm* and loading the file into BASIC:

```
LOAD "TRAXX",8,1 RUN
```

Alternatively you can assign a directory or disk image to a drive number (device 8 to 11) using the menus. As on the C64, which used the same serial peripheral interface and drives, other device numbers are used to communicate with printers, serial ports and other peripherals.

*cbmconvert* compiles readily on Linux if you unpack the archive and type *make unix*. *cbmconvert* can handle disk images and many other formats, but gets confused by upper-case file names and extensions in programs compiled on lesser systems, so you may need to rename input files.

Another attempt to cut costs brought forth the obscure C116, a re-packaged C16 with the typewriter keyboard supplanted by wobbly calculator-style keys. A VIC-10, based on the VIC20 hardware and PET 2001 keyboard, surfaced in Japan.

## Plus 4

Plus 4s have 64K memory like C64s, and the extra ROM of the C16, all packaged in a slimline black box. The name comes not from Edwardian trousers but the four productivity programs (rather feeble 'business' software) in an extra 32K of bank-switched ROM, first known as Three Plus One.

## XPET and XCBM2

**X**PET emulates several models of the all-in-one PET, the *VICE* emulator options let you switch through the early years of PET development, with support for the buggy 'old ROM' and tape files as well as disk images.

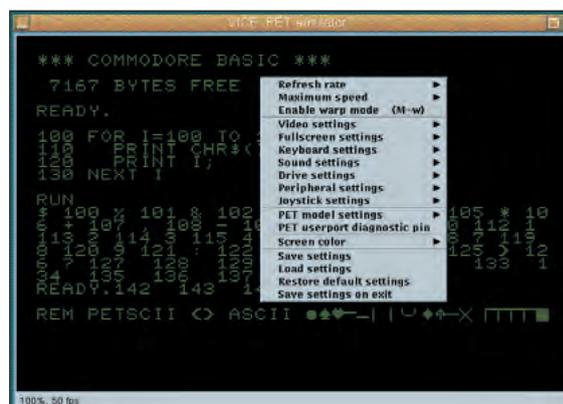
The CBM2 computer was also known as the C610. Again *VICE* emulates several models from the range, though not the later multi-processor versions. These are authentic emulators, but the firmware is basic in every sense, and mono character graphics limit applications. Unless you've used the originals, *XPET* and *XCBM2* are likely to leave you cold.

## PfauZeh

*VICE* competes for the VIC-20 emulation crown with *PfauZeh*, which puns on German pronunciations of the letters V and C. *PfauZeh* stalled in 1998 at version 0.8 on Linux, though work on the Win32 version continues. It is supplied as over a megabyte of compiled binary, with ROMs, HTML files and an example game. *PfauZeh* ran on my Debian 2.2 installation once I reset the video depth to eight bit to provide access to palette hardware, and switched my usual *Blackbox* window manager for the minimalist *TWM*, freeing up unallocated colour for the emulator.

*PfauZeh* initially presents a titchy X window, though two pixels are allocated to every one on a VIC line to keep the aspect ratio about right. The *-hires* command line option gives a larger view, plotting four pixels in X and two in Y for each original one.

Screen updates default to 30Hz – one for every two fields



*XPET*, showing a *VICE* menu and PETSCII characters.

drawn by the original hardware – but you can crank this up to 60Hz for the default NTSC video mode, emulating the original 6560 VIC chip, and from 25 to 50 Hz for the 6561 PAL version. Three menu options select the video format, frame rate from 1:1 to one in six, and nine or 28 lines of overscan at the top of the screen. Sliders pick the original 5K RAM, or 3K, 8K, 16K or 24K expansion, with a toggle for extra RAM at address \$A000.

A delay knob slows emulation by wasting from zero to 500 cycles every time an instruction is fetched, but *PfauZeh* is hardly speedy. A 10,000 iteration FOR loop runs across two BASIC lines in about 16 seconds on my K6 with 60Hz display, and in 12 seconds at 30Hertz, about as slow as the original. ➤➤

## On the web

This series has reported on more than 100 emulators already, and the results now nestle at <http://simon.mooli.org.uk/LXF>.

So far this site only contains early parts of the series, but it will eventually host the lot, plus material that didn't make it into the magazine.

# Emulators

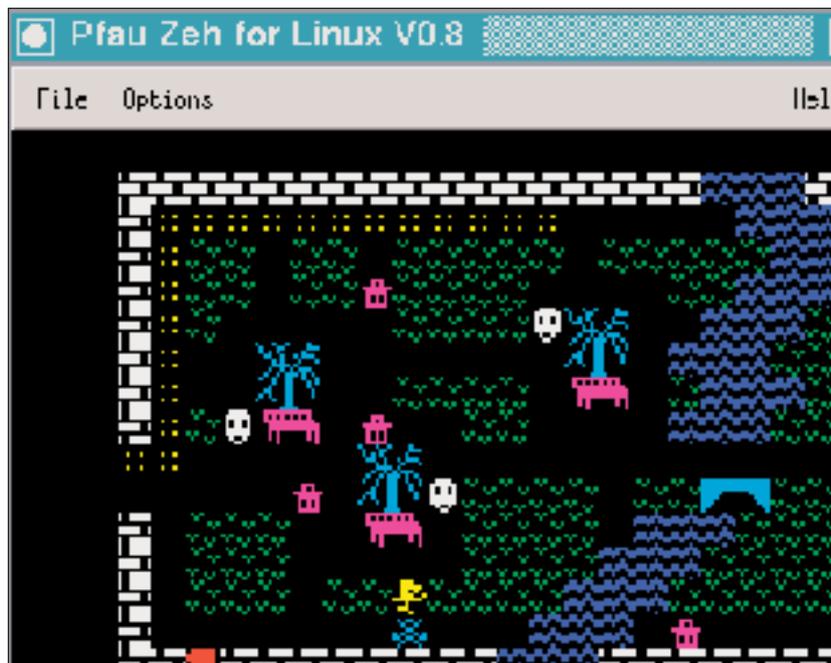
« A file menu will load or save BASIC and plug in or remove a ROM cartridge, with reset implied by removal. Another menu option saves display, CPU and memory configuration settings. The **-joy** emulation maps a Linux joystick to mimic a VIC one. The alternative keyboard joystick emulation uses the right **Ctrl** and cursor keys, or numeric pad arrows with **NumLock**, for directional moves. Two analogue paddles are emulated with a mix of cursor and numeric pad keys. The mouse can also simulate a lightpen, for those few VIC programs that expect one. The mouse pointer shows where the pen is pointing and any mouse button emulates the one on the pen.

*PfauZeh* emulation seems precise, with undocumented 6502 opcodes, interrupts and instruction timings correctly modelled, plus both VIA interface chips and all VIC graphics modes, including the 16 colours, screen position and raster beam registers. Video output is line based, rather than synchronised with whole fields or individual pixels, so it copes with programs that reprogram the VIC in horizontal flyback but not those few that rely on precise timings to tweak it during a line.

The HTML documentation is a good read, neatly but rather garishly written mainly in red or yellow on black. It contains no links and only one graphic, showing the original VIC keys and their PC equivalents. The author lists many programs known to be compatible, with advice on memory size, PAL and NTSC dependencies, and **SYS** commands to get programs running. **SYS** invokes a block of 6502 code on many Commodore machines by diverting the program counter to a given address.



**Commodore's PET 2001 was the first all-in one personal computer, but the fiddly key grid was bad news for touch-typists.**



***PfauZeh's Thriller* demo is mostly written in BASIC.**

## Compatibility

If your files are intact *PfauZeh* should run almost anything you throw at it. Only two programs fail outright, though half a dozen others show minor timing or graphical glitches and one of the ones that doesn't work is OK in its NTSC incarnation and only bombs on PAL. Sound supports the eighties standard of three tones and one noise channel, converted to eight bit mono at 44.1 kHz, though I found this noticeably jerky on the games I tried, whatever the update options. You can run *PfauZeh* without sound.

Keyboard input is also erratic even on a 500 MHz K6, with

momentary hiccups in the display update as every few characters are entered into BASIC. The **F11** key or a menu option reset the VIC, **Esc** quits and **F2** and **F4** are shortcuts to BASIC **LOAD** and **SAVE** requesters, which use *Motif* file selection. I was not able to get save to work, despite checking file permissions.

The tab key stands in for Commodore's **RUN/STOP** button, and left **Alt** for the special **CBM** key. **Insert** generates a £ or Lire sign. Cursor movement is authentically but irritatingly achieved using **PageDown** and **PageUp** to move down and right respectively, with **Shift** to reverse the direction. **F12** acts like the VIC's **RESTORE** key and the letters, numbers, space bar, **Enter** and **Home** keys work as expected. Odd numbered function keys from **F1** to **F7** mimic the four original VIC function keys, with even numbered keys assigned to adjust the emulators frame rate or memory size.

Jerky sound and keyboard input spoil what is otherwise a comprehensive emulator, especially in its graphics. More refined host device access could make it attractive, and while Arne Bockholdt is no longer involved with Linux he is willing to help others who'd like to update it.

## XVIC

*XVIC* is part of the *VICE* package, like *XPET* and *X64*. Like those it shows the frame rate and percentage of original CPU time in the bottom window border. It defaults to 50Hz and authentic 100% CPU speed but can manage over 400% and 200 frames per second if you take off the brakes, and ten times the original CPU speed at about half the authentic frame rate if you slew the options that way. On the same 8-bit X11 window with two for one pixel mapping and the *XVIC* default 50Hz display update it was as fast as *PfauZeh* with its half-NTSC 30Hz video, without the sound and keyboard hiccups, and four to eight times faster than the real thing.

*XVIC* benefits from *VICE* speed-ups like Warp mode, video caching and configurable sound sample rate. X display depths may range from 1 to 32 bits per pixel, with four drives emulated by **LOAD** and **SAVE** commands from BASIC, rather than menus that bypass the VIC operating system.

Speed limits are by percentage of full speed, and may far exceed 100 on a modern Linux box. Double size, double scan and full screen modes are controlled from the command line or two full sets of menus accessed by left and right mouse buttons.

Four RS232 devices with individual baud rate settings, normally connected to `'dev/ttyS0'` and `'dev/ttyS1'`, the Linux serial ports, the file `'rs232.dump'`, and a pipe to `'lpr'` for printing. Access to the VIC user port can also be diverted to a *VICE* RS-232 emulation, and four printer devices are similarly linked to files, Unix devices, or *VICE* IEC port emulation.

1541, 1571, 1581, 2031 or SpeedDOS disk drives can be emulated down to the level of the drive firmware, which is very accurate but costs a lot of time, or implemented by ROM patches that fool most software at lower overhead, and mapped to host directories either way.

Two joysticks and the 1351 proportional mouse can also be emulated by Linux peripherals. *XVIC* has impressive sound, well used in the Jeff Minter games on our cover disc, but struggles with some of the Llama-man's graphical jiggery-pokery, especially on menu screens where you may just get erratic flashing rather than Jeff's trademark beam synchronised psychedelia.

## C16 and Plus 4

The only C16 emulator I found specifically for Linux was *xc16emu* version 0.5.4, written in 1998 for Qt 1.3 or 1.4, and I had to comment out some cosmetics to fit the C++ to the Qt 2.3.0 on my Debian box. I replaced the caption in line 71 of *xc16emu.cc* with a literal string, removing lines 53 and 54 in *mainwindow.cc* and the body of the **about** and **info** methods there. The result is on your coverdisc.

Video and CPU ran at about double speed on my 500MHz K6 and Matrox G400, using the default 1:1 pixel window. Scaling the screen up to Zoom 3 or Zoom 4 brought authentic speed and full-screen. *xc16emu* has been tested on MIPS and Alpha.

Load the four simple demos in the BASIC subdirectory with **DLOAD "BASIC/file.BAS"**, press TAB to stop them. Once tweaked, with ROM images from funet, *xc16emu* offers a pretty GUI, let down by imprecise CPU timings and lack of sound.

### Plus 4 emulators

The *Minus4Java* applet emulates the Commodore Plus 4 and C16 slowly but well enough to run many games, redirecting cursor key-presses to the emulated joystick. The implementation as an Applet means that it cannot save any results. A 500MHz Celeron, Netscape *JRE1.2*, RH 7.0 ran at 50Hz window updates.

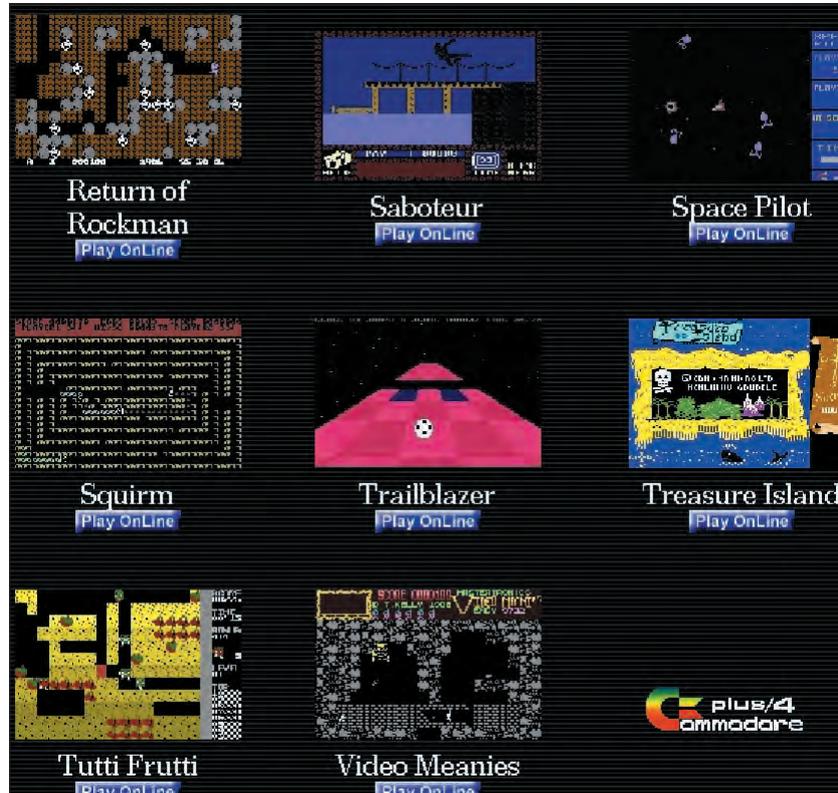
### YAPE

*Yet Another Plus/4 Emulator*, written by LXF reader Attila Grosz. Attila lost his Mandrake set-up recently and has been awaiting an upgrade before re-starting work. The GPL'd *YAPE* is written in C++, and requires *LibSDL 1.2.1* or later.

*YAPE* compiled without demur on my Debian 2.2 test system, with *SDL 1.2.3* installed. Just type **make** then **./yape** to get a familiar Commodore BASIC screen in an X11 window. *SDL* console messages warned it was unable to find the desired 8-bit unsigned mono audio output at 44,100Hz, but settled for a close equivalent mode supported by my SoundBlaster Live. After the command you can specify the name of a TAP tape image or a PRG file holding program bytes ready to be loaded directly to memory. Files load from the '~/.yape' subdirectory.

Once *YAPE* is running, the first few function keys give quick access to BASIC commands like **GRAPHICS**, **DLOAD**, **DIRECTORY** and **HELP**, as on a real Plus 4. **F5** and **F6** simulate cassette drive **PLAY** and **STOP** keys. **F7** saves screenshots in the bloated BMP format – *SDL* makes some things too easy! The start screen comes out as a 143,350 byte BMP, shrinking to 1819 bytes once *GIMP* converts it to PNG. **F10** saves current settings to your home directory. **F11** with optional **Ctrl** or **Shift** triggers three types of +4 reset, and **F12** and **Esc** quit.

Left **Alt** with **G**, **J**, **I** and **S** toggle a display of the frame rate above the main window, joystick emulation with Linux cursor keys,



the emulated joystick port, and a frame rate limiter, clamped to 50Hz rather than 50MHz as mistakenly documented!

Sample Plus 4 games online in *Minus4* at [plus4.emucamp.com](http://plus4.emucamp.com)

### Key points

*YAPE* keyboard handling is conceded to be a bit shaky, and depends on a table mapping *SDL* codes to Plus 4 matrix positions. **End** sends the £ or Lire key code, but I found the zero keys did not work at all with my US PC keyboard. After some fiddling with the source and tentative key-pressing I found out that **0** is *SDL* key **\$30** and that the key to the left of **1**, with a back tick and tilde on it, generated the zero code, which is **\$60** to *SDL* (and ASCII) and 28 for the +4. That key is not part of the 64 key Plus 4 matrix, so the emulator is usable once you know to press that when you want a zero. Arguably this is more logical than the normal arrangement of having 0 after 9, but like most people I've got used to the latter scheme, so I changed the first entry on line 29 of keyboard.cpp from 38 to 28 and got the normal zero key working, as well as the one before **1**. Even with a weird keyboard you should be able to remap the emulated keys this way. Tweaked source and a compiled version of the emulator is on your coverdisc with the author's original archive. [LXF](#)

## Links

A web of Commodore sites:

C16 Deutschland: [www.c16.de](http://www.c16.de)

C for CBMs: [www.acc.umu.se/~arvid/cc65\\_mirror](http://www.acc.umu.se/~arvid/cc65_mirror)

Link cable: [www.lb.shuttle.de/puffin/cbm4linux](http://www.lb.shuttle.de/puffin/cbm4linux)

LXF emus online: <http://simon.mooli.org.uk/LXF>

Minus4: [www.javalemmings.com/minus4/files](http://www.javalemmings.com/minus4/files)

PfauZeh home: [www.classicgaming.com/pfauzeh](http://www.classicgaming.com/pfauzeh)

Plus 4 online: [www.plus4.org](http://www.plus4.org)

VIC-20 ammunition: <ftp://ftp.funet.fi/pub/cbm/vic20>

VICE home: <http://viceteam.bei.t-online.de>

XC16: <http://ibiblio.org/pub/Linux/system/emulators/>

YAPE: [www.keepitretro.com/plus4](http://www.keepitretro.com/plus4)

# Tutorials

Our experts offer help and opinions on a whole host of Linux applications

## Your guide to getting things done!

Whether you are just starting out in Linux, or an experienced veteran, there's always more to learn. Every issue of *Linux Format* is packed full of practical advice, and nowhere is it more concentrated than in our tutorials pages.

Here you'll find expert guides to all sorts of things, from Basic Linux usage to understanding and deploying network solutions, from simple script coding to the complexities of Perl regular expressions, Java server apps and more. We aim to bring a good mix of tutorials to each issue, but if you have any suggestions for topics you'd like us to cover, why not contact us, by post, by email ([linuxformat@futurenet.co.uk](mailto:linuxformat@futurenet.co.uk)) or log on to our web site and post your suggestions in our special forums? ([www.linuxformat.co.uk](http://www.linuxformat.co.uk)). Hope to hear from you soon!

**Nick Veitch** EDITOR

## How code is represented

Including code in magazines can be tricky, but we hope our notation will help it become clear. When lines are too long for our columns, the remaining text appears on the next line in a solid blue box:

```
procedure
TfrmTextEditor.mniWordWrapClick
(Sender: TObject);
otherwise, there is usually a gap
between lines:
begin
mniWordWrap.Checked := not
end;
```

Usually, you'll find the code on our CD/DVD too.

## THIS MONTH...

### File permissions >>

Are you confused by Linux's access control system? Our guide will help you sort out your users and groups. **p68**

### Apache

Get to grips with dynamic web sites as we look at setting up CGI and PHP in this next part of our guide to the popular web serve. **p72**

### Perl

Charlie Stross is always on the look out for new uses for Perl. This month: how to play your MP3 collection. **p76**

### Java

Parsing XML is easy, but how do you make use of that data? The next part in the series looks at how to create Java objects from XML files. **p80**



### Kylix

We look at the lighter side of programming with Kylix. This time, breathe some life into your applications and find out all about doing animation. **p82**

### CGI

To build an e-commerce site, you'll want to use sessions to manage your user's data. Find out how in the last part of this server-side series. **p86**

## TIP OF THE MONTH!

**Whether you are a home user or a sysadmin**, there must have been some point when you found it useful to look at the system logs. We could do a whole feature on system logs, and indeed we will, but here are some useful tips.

First of all, you should take a look at what logs are available and where the messages are going. You could just look in the directory `/var/log` and work out what process logs output to where, but it's much more informative to look

## Tailing logs

at the file `/etc/syslog.conf`, which will actually tell you what warning levels are being logged from where and into which files.

The second useful thing to know is the tail command. This is a text viewer, which simply outputs the last ten lines of a file. So, for example, `tail /var/log/messages` will tell you what's been going on recently on your system.

If ten lines is too few or too many, you can change the number output by

tail with the `-n` switch. For example

```
tail -n 20 /var/log/apache/access.log
```

Even more useful for debugging processes and daemons is tail's `-f` switch, which 'follows' the file. This makes tail poll the file continuously at set intervals, and refresh its output as information is appended to the file.

```
tail -f /var/log/daemons/errors
```

You need to open this in another console window, but it will give you a running account of the information added to the file.



```
-rw-r--r-- 1 mkelly users 22 Dec 6 22:21 sample2.txt
```

This sets read access for all users, but leaves all other permissions alone. It is also possible to apply an individual permission in this way. Say you wanted to grant execute permission on all your files for the owner, but didn't want to grant to anyone else. This is executed as follows:

```
$ chmod u+x sample2.txt
```

```
$ ls -l sample2.txt
```

```
-rwxr--r-- 1 mkelly users 22 Dec 6 22:21 sample2.txt
```

So as you can see it is possible to allocate permissions absolutely using an octal string, or to simply update individual permissions. There is more to the `chmod` command than we can cover here, but have a look at **man (1) chmod** for more info.

## The Root User

So that's mortals covered, what about the all-powerful super-user – how does root fit into this scheme. Put simply, root has access to every file on the system, which is why it is so important to log in as a normal user. Root can delete/mangle any files regardless of permissions, and using root as your day to day user account simply increases the possibility of you doing something foolish. File permissions apply to files belonging to root in the same way as for normal users. If root decrees that you can read/write/execute a file belong to him/her (either by allowing access to a group or to "others") then you can alter/use that file, in the same way that you can access one of your fellow users' files with the same permissions.

Whilst you can override any permissions with root, you have to remember not to go insane. Linux may seem like it has a very strange directory structure, but it is one that has evolved over time. Some file locations are very important, so it's not wise to go moving stuff to new homes at random. Likewise, while it may be tempting to save everything in the root directory just remember how much easier it would be to delete everything on your system with a slip of the `rm` command while in `.` While not every distribution places files in the same place, it is worth a read of the Filesystem Hierarchy Standard at [www.pathname.com/fhs/](http://www.pathname.com/fhs/).

## File Ownership

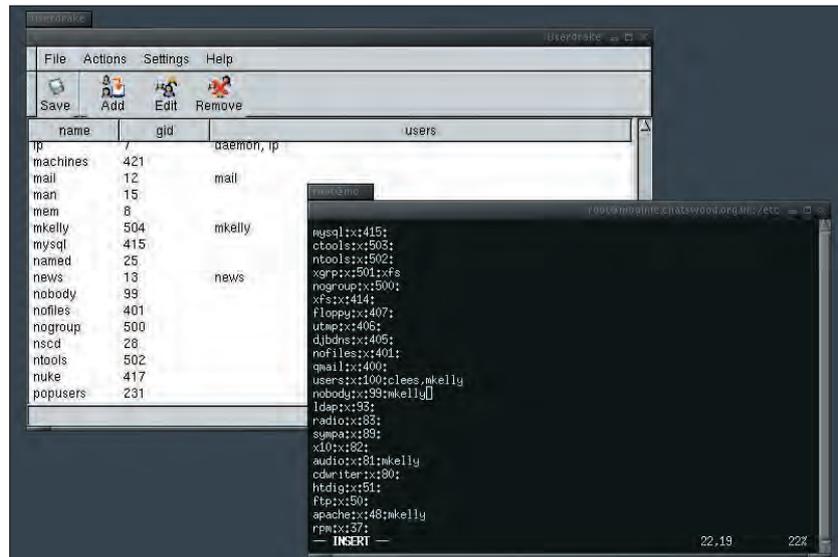
While discussing permissions above the concept of file ownership was introduced. All files on your system will belong to someone and many of them will belong to root. When you create a file it will automatically belong to you, and usually your default group. File ownership does not usually change that much as users are tightly restricted in what they can do. Much ownership changing will be done by the root user as he/she can decide the ownership of any file on the system.

The lack of freedom when it comes to changing ownership is not without good reason. If you could change file ownership to any other user it's possible that you could exploit otherwise safe situations creating a security risk on your system. Fortunately more use can be made of the group attribute, which is more lenient and allows you to assign a file that you own to any group you are a member of. We will use the `chgrp` command to assign our files to other groups in the example below.

Detailed information on both of these commands is available in the man pages **man 1 chown** and **man 1 chgrp**

## Protecting Your Web Docs

Say your Linux box is a web server on a local network or even on the Internet. Your own personal html directory contains sensitive



information that should only be accessed by certain people through password-protected pages. Maybe it's not that tight, but you've got a bunch of scripts that you don't want others to see the code behind. You can prevent prying eyes viewing from the web, but what about local users? Proper use of permissions on your own html directory will protect your pages against marauding local users. Of course you need to be careful about how the permissions are set up. You don't want to deny access to the person who maintains the info, but you certainly don't want to be so restrictive as to prevent the `httpd` daemon from being able to read the files in order to serve them up to a web browser.

On a standard Linux box running `Apache`, the personal html directory is called "public\_html" and is located in your home directory. When `Apache` receives a request for `http://localhost/~username/` it fetches from the directory `/home/username/public_html` (or `~/public_html` for short.) This may not be set up on your system in which case you'll need to do so yourself, or get your administrator to do so.

My `Apache` `httpd` daemon runs under the user "nobody" and group "nobody". In the remainder of this example where you see `~/public_html` you should substitute your own personal html

**You can add yourself to more groups by editing the `/etc/group` file, or using a graphical tool supplied by your distro – the choice is yours.**

## Default Permissions

### Behind the mask

This talk of permissions begs the question "Do I have to continually use `chmod` every time I create a new file?" On most systems, non-executable files are created with **644** permissions. This of course gives read access to all your files to everyone. If you want to avoid this, say by making new files appear as **640**, you need to use the shell built-in command `umask`. This is generally set system wide in `/etc/profile` (for `Bash` users) but you can override it in your own `~/.bash_profile`.

It works by specifying a mask for the creation of new files – you subtract the mask from the value **666** in order to get the default permission of new files. For example, in order to get a permission of **640**, you need to subtract **640** from **666**, which results in **026**. Okay, it's actually not quite that simple – that covers the

creation of ordinary files. When an executable gets created it is going to be made executable by everyone by default, so you should add 1 to the mask for every category of users you don't wish to be able to execute the file. So if you had a script or program that you wanted people in your group to be able to read but not execute, and deny everyone else access completely, you will need to add a 1 to the group permission mask and the other users permission mask, making your complete `umask 037`.

You can set your own mask (either temporarily at the command line or permanently in your start-up file) by issuing `umask mask_value` (where `mask_value` is something like **037**.) You can view the current setting by issuing `umask` at the prompt, or `umask -S` for a more explanatory notation.

# LinuxFormatTutorialFilePermissions

« directory, and where you see the username “nobody” and groupname “nobody” you should use your own equivalents. You can check what user/group your copy of *Apache* runs under by checking your *httpd.conf* file. Where you see the username “mkelly” you should substitute your own username.

On our system we want the user “mkelly” to be able to read and write files within `~/public_html` and its sub-directories, but all other users should be denied access (both read/write and execute.) The files are currently readable by anyone, and so the first thing we want to do is clamp down on who can read/write those files. To turn off read/write/execute permission on the `~/public_html` directory we must enter the command:

```
$ chmod o-rwx /home/mkelly/public_html
$ ls -l
total 1
drwxrwx--- 18 mkelly users 1024 Apr 9 13:25 public_html
```

Perform an `ls -l` as shown above, and note the permissions only extend to the owner and the group. Change into `~/public_html` and repeat there and in any sub-directories in order to protect the files below. Now only user “mkelly”, and anyone in the group “users” can read or write to the web documents. Unfortunately, while this is a step to what we want, it’s partially a step backwards. If you fire up your web browser, and point it to “`http://localhost/~mkelly/`” (substituting your username for “mkelly”) you will likely get a **403 Forbidden** error. To top it all off, most of the other users are still able to read our files in the actual directory as they are members of the group “users”.

Fortunately, fixing both these problems can be done with a single command. The `chgrp` command can be used to specify which group a file belongs to. First though we need to add our user to the group “nobody”. As root, open the file “`/etc/group`”

Right: **Konqueror is the file manager for KDE2 and offers the user an easy way to manage permissions for a file.**



and find the line for the group “nobody”. Add the desired username to the end of the line (see the screenshot for an example) and save the file. Now change the group of `~/public_html` using the `chgrp` command (or your favourite file manager software.)

*Voilà!* If we repeat that in sub-directories of `~/public_html` the *httpd* daemon should have permission to read those documents, and other normal users should not be able to read, write or execute those files. Attempting to access those confidential documents as another user should result in something similar to:

```
$ cd /home/mkelly/public_html
bash: /home/mkelly/public_html: Permission denied
```

This was quite useful but there is something you should watch out for. When we create new files we need to set the proper permissions and ownership afterwards. On most systems creating a new file will cause it to be owned by the user, in this case “mkelly”, and belong to the group “users”. Most systems also create files with the permissions `rw-r--r--`. For example, if I create the file “`test.html`” in `~/public_html`, I can see that it gets created as follows:

```
$ ls -l test.html
-rw-r--r-- 1 mkelly users 81 May 4 22:28 test.html
```

We need to do two things here. We want to get rid of any permissions for others and make it readable to the *httpd* daemon, so we do

```
$ chmod o-rwx test.html
$ chgrp nobody test.html
```

See the box *Default Permissions* for details on how to generate new files that don’t have to be altered using `chmod`.

## Group File Sharing

Another use for permissions is file sharing for a group of people. An example situation is that of a shared project where a set of common files are used by the project members. Granting read/write permission to “others” would allow everyone to use the files, but offers no security. A better idea is to create a group (e.g. “`projgrp`”) and assign all the files to that group.

Creating a group should be done by root. From the command prompt issue the `groupadd` command with the name of the new group, for example

## Changing Permissions In KDE/GNOME

### The way of the GUI

Permissions don’t have to be modified from the command line, you can use your GUI as well. In GNOME start up your preferred file manager – I use *Nautilus*. Navigate to the file you wish to change permissions on, and right click on it. From the menu choose the “Show Properties” option. A new window will open displaying the file properties – click the third tab (marked “Permissions”) and you’ll be presented with something similar to what you see in the screen shot.

You should see a graphical representation of the current permissions on the file – clicking on the tick boxes will cause that permission to be set, and *vice versa*. (The options Set UID, Set GID and Sticky will be explained in the section “Special Permissions” below.) The “Text View:” and “Number View:” lines should update as you change options – observing this is a good way to familiarise yourself with common modes.

KDE users should start up *Konqueror* and navigate to their chosen file. As with *Nautilus* you right click on the file and select the “Properties...” option. The tab marked “Permissions” displays an almost identical arrangement to that produced by *Nautilus* without the number and text representations of the permissions. As before setting a permission is as simple as clicking on one of the buttons. (See the screenshot for a view of the dialog)

The permissions tabs in both dialog boxes also showed the Ownership section. Note that as described in the section “File Ownership” you can only change group ownership to a group of which you are a member. Until you are a member of multiple groups you won’t be able to change the group to anything else.

**# groupadd projgrp**

Your new group will be created. To add each of your prospective group members we could edit the file "/etc/group" as we did before, but we will use command "usermod" as follows:

```
# usermod -G [existing group1,existing group2,...,]projgrp
username
```

Take care in issuing this command – you must specify the user's existing groups otherwise he/she will be removed from them. For example if the user was already in the groups "ftpuser" and "pppuser", then you would issue **usermod -G ftpuser,pppuser,projgrp username**. You can find out what groups a user is already in using the *id* command at the prompt.

Use the *chgrp* command to change the group of all the files to your new project group. Each member of the group who wishes to contribute files must do this. Now if you restrict file access to user and group only (**chmod o-rwx filenames**) you will find that only the members of your project group can actually edit, view or execute your files.

While groups may seem useless on a single-user system bear in mind that many processes and daemons on your system run as their own user and group (like the *httpd* daemon we saw above), and so assigning files to groups is an added layer of protection for those files without them belonging to root.

## Special Permissions

We didn't quite give you the full story when we said that the permissions were read, write and execute. There are in fact three other special attributes used in relation to executable files – *setuid*, *setgid* and the *sticky* bit. The *setuid* bit will cause the process executed by the file to run with the owner ID of the file (useful for users who need to use a program that needs root access to a device – e.g. *pppd* for modem dialup use.) *Setgid* is similar to *setuid* but sets the process to execute with the group ID the same as that of the owner. The *sticky* bit (save text bit) will save the image of the program in the swap memory so that it can be executed more quickly in future (does not need to be reloaded into swap.)

## Directories

Directories have permissions too, but are not quite the same as files. Directories use the read/write/execute bits as well as the *setgid* and *sticky* bits (but not *setuid*.) A directory with read access will allow listing of directory contents. A directory with



**Nautilus is the GNOME equivalent of Konqueror and works in much the same way - permission editing in a file manager is pretty similar no matter what you use.**

## Usedir

### Setting up your ~/public\_html on Apache

For this exercise you want to be able to use the **UserDir** feature of *Apache*. To do this, as root, open up your *Apache* configuration file called "httpd.conf" (or in "srm.conf" on older systems.) Search down the file until you see a configuration directive called **UserDir**. It should read

```
UserDir public_html
(you may simply have to uncomment it.)
```

This enables the **UserDir** feature, but you need to tell *Apache* that it's okay to read

documents from there. Below the **UserDir** directive enter the following text

```
<Directory /home/~username/>
Options Indexes
</Directory>
```

This should be sufficient to allow you to follow the tutorial. Now restart *Apache* to allow the changes to take effect, and try your new html directory out by loading <http://localhost/~username/> into your web browser.



write access can have its files added, removed or altered. A directory with execute permission can have its file details listed (e.g. in **ls -l**) Therefore having read but not execute permission causes you to be able to look at directory contents, but not see full file details. Note that execute but not read permission does not allow any listing of files or file details.

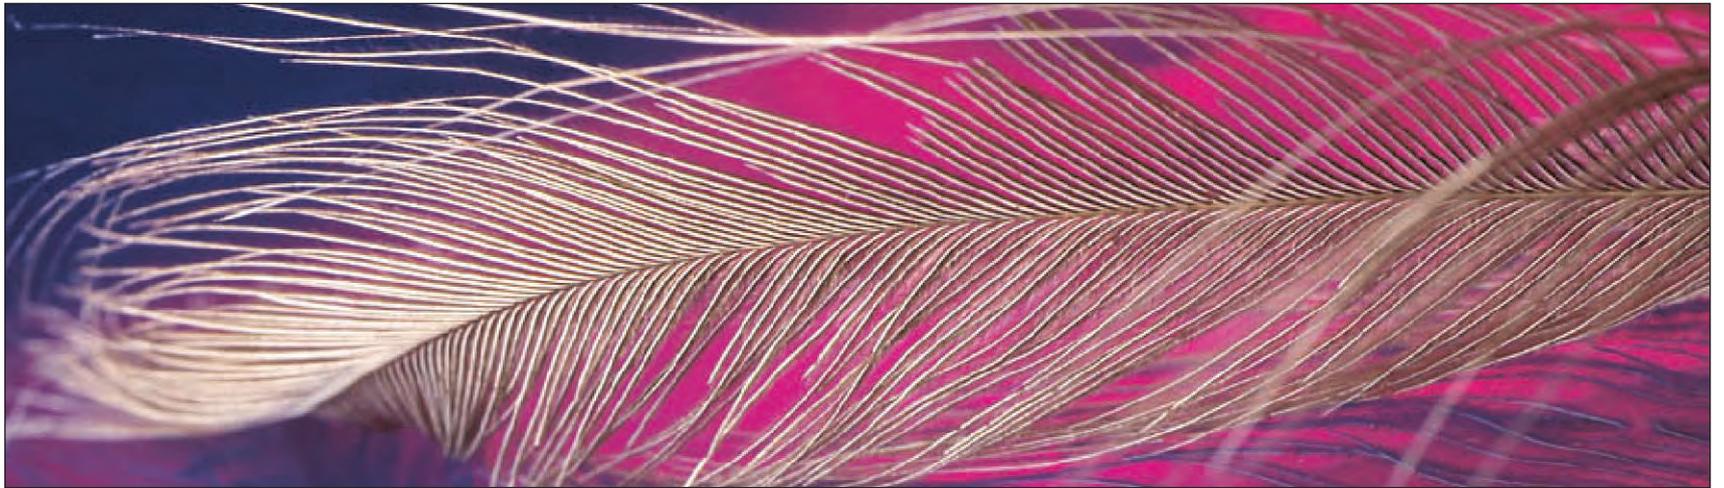
Using *setgid* on a directory will cause files created in a group to take on the same group ID as the directory instead of the group ID of the file creator – this is obviously useful when creating files in your *public\_html* or project group directories in the previous examples. Using the *sticky* bit makes a directory become "append only" – users cannot remove files unless they own them.

In the example above regarding maintaining security of web documents, you may remember that we denied read/write access to all web documents to any users except those we specified as our own user and the members of the group "nobody". However, it is still possible for a local user to look at the directory contents, so we may want to protect those contents from onlookers (at times filenames can give away as much information as file contents.) To protect your web documents, you can set the permission **rwXrwx---** on your *~/public\_html* directory – that way only the owner and members of the group "nobody" can read or write to the directory – this also applies to sub-directories. Now your web-documents are safe from both tampering and viewing by unauthorised local users.

## Finally...

Permissions can be a pain at times – slowing you down when you just want to get the job done quickly. But the next time you do an inadvertent **rm**, and a file permission saves your documents, you may just be glad you protected those files after all! [LXF](#)

**The result of the ~/public\_html directory being accessible to user "nobody" – we're now officially Web-published!**



## DYNAMIC CONTENT

# Apache webserver

**PART 4** Chris Brown looks at the world of server-side scripting, and configures Apache to produce nicely formatted directory listings.

**W**hat happens if you enter a URL into a browser which names a directory, rather than a file, on the server – e.g. by entering a URL such as `http://www.apache.org/`? Effectively, this is a request to retrieve the top level directory (*i.e.* the **DocumentRoot**).

Apache's default behaviour is to deliver the file `index.html`, if it exists. You can specify different behaviour with **DirectoryIndex**:

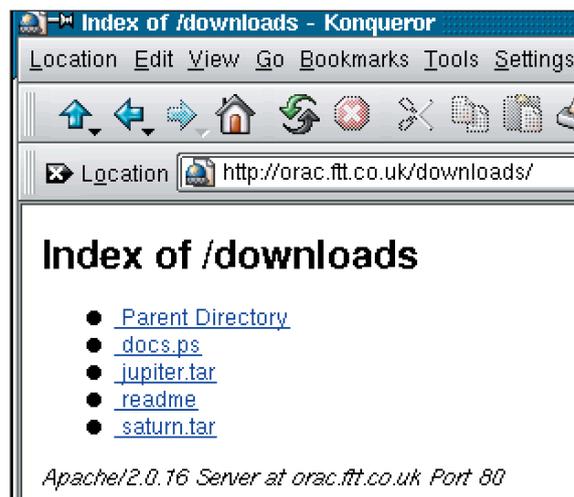
```
DirectoryIndex home.php index.html greeting.txt
```

Apache will look for each of the named files, in turn, and deliver the first one it finds. If you put this directive up at the top level of `httpd.conf`, it will apply to all directories. If you put it inside a container directive for a specific dir, it will apply only to that dir.

Alternatively, Apache can generate a listing of the directory for sites that provide files for download. Let's create a directory called `downloads` under **DocumentRoot**, and put some files there:

```
# cd /usr/local/apache2/htdocs
# mkdir downloads
# cd downloads
# touch readme jupiter.tar saturn.tar docs.ps
```

Fig 1. Plain directory listing.



`touch` just creates empty files. For this demonstration, we don't care about their contents, just the fact that they are there.

With these files in place, accessing the URL `http://localhost/downloads` results in the browser listing shown in **fig. 1**. Apache is internally generating the HTML code which results in the browser display shown. The file names are links you can click on.

## ...and fancy

We can spice up our directory listing with the **IndexOptions** directive. Firstly, we can turn on **Fancy Indexing** in `httpd.conf`:

```
<Directory /usr/local/apache2/htdocs/downloads>
IndexOptions FancyIndexing
</Directory>
```

Here, we have placed the **IndexOptions** directive inside a **<Directory>** container so that its effect will be restricted to the `downloads` directory. In the browser, the directory listing now looks like **fig. 2**, with additional columns to show the file sizes, etc. The column headings are themselves links; clicking on them re-submits with a request to re-sort the display on that column.

Let's go further and add icons and text descriptions which give info about file types, based on filename extensions. Apache includes a dir called 'icons' (under **DocumentRoot**) containing about 70 small icons as `.gif` files. You can associate icons with file types using the **AddIcon** directive, as shown below. You can also specify a default icon to be used for unrecognised file types, using the **DefaultIcon** directive. Finally, you can associate descriptive text with a file type using the **AddDescription** directive. Putting all these together gives us the following entries in our `httpd.conf` file:

```
<Directory /usr/local/apache2/htdocs/downloads>
IndexOptions FancyIndexing
AddIcon /icons/tar.gif .tar
AddIcon /icons/a.gif .ps .ai .eps
AddDescription "Tar archive" .tar
DefaultIcon /icons/unknown.gif
</Directory>
```

The resulting directory listing is shown in **fig. 3**. The **IndexOptions** directive has numerous options to fine tune the directory listing even further. You can make the icons part of the clickable link for the file name; suppress specific columns of the listing; control the width of the 'Name' column; and so on.

## Server Side Scripting

We've covered a lot of ground in these tutorials, but all our efforts so far have resulted in *Apache* delivering static HTML content to the browser: the content remains the same, unless someone goes in and changes the HTML files. The kinds of apps run on the web these days require content delivered to the browser to be generated dynamically at the server end, e.g. by taking data from a database and formatting it as HTML. The app may also capture and process items on a web form, update a database and connect with some other service. For these apps we need some way in which the web server can run application code.

This technique is generally called server-side scripting, and quite a few technologies have emerged to support it. We're going to take a look at two rather different approaches to server-side scripting, popular with *Apache* – CGI scripting, and PHP.

## Creating CGI Scripts

At its simplest, a CGI script is a program launched (as a separate process) by the web server whose output is captured by the server and returned to the browser. The server doesn't examine or modify the script's output in any way, so the script had better generate something that the browser can make sense of! Calling these programs 'scripts' is really too narrow a title. They can be fully-compiled programs written in e.g. C, or in interpreted (i.e. scripted) languages. Perl is popular because it excels at the kind of data handling required in CGI programs. However, to avoid having to learn another programming language we'll use a shell script for our examples. **Fig. 4** shows the basic CGI architecture. Notice that the web server generates the HTTP header for the response, but the CGI program generates the content.

First, we'll create a simple script which displays the time and date (courtesy of the **date** command). To conform to the rules of CGI, the script is also required to generate a one line HTTP response header to tell the browser that the context is plain text, followed by a blank line (which tells the browser where the end of the HTTP header is). Here's the script, which we'll call 'dater':

```
#!/bin/sh
echo Content-type: text/html
echo
echo -n "Date and time are : "
date
```

The first line of the script specifies which interpreter is to be invoked (in this case the shell) when the script is executed. If the script was written in Perl, this line would instead be:

```
#!/usr/bin/perl
```

Now come a couple of tricky questions. What permissions should the file have? Who should own it? Where should we put it? To answer these we need to establish the scenario in which the script is to be used. We'll resurrect our user 'minx' who appeared in a previous article, and assume that the script is to be owned and maintained by minx. So the file should be owned by minx and have **rwX** permission for its owner, so he can modify and test it. Remember, though, that when the script is invoked by *Apache* it will be started by a process running with the user and group identities defined by the User and Group directives in the httpd.conf file. In our previous examples, we defined a user called 'webuser' and a group (also called webuser) for this purpose. So to ensure that *Apache* can execute the script, we'll change its group ownership to webuser, and turn on group execute permission, like this:

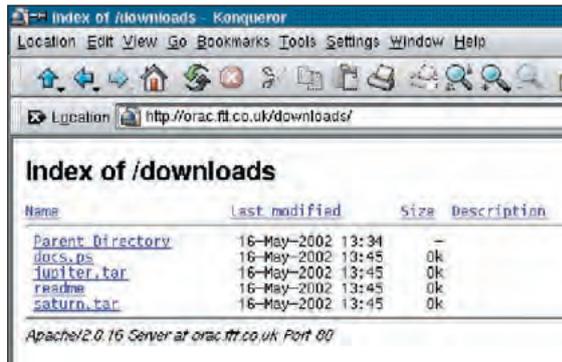


Fig 2. Fancy directory listing.

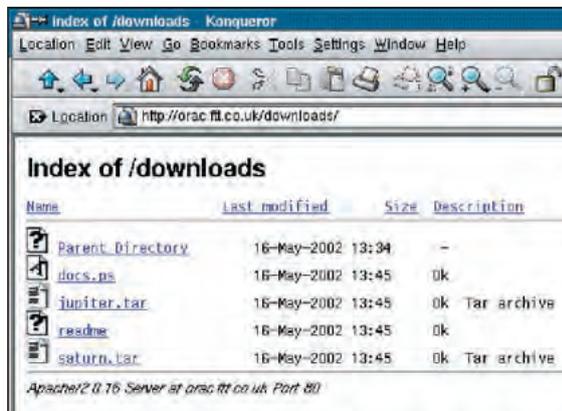


Fig 3. Directory listing with icons.

```
chown webuser dater
```

```
chmod ug+x dater
```

(Done as root, because only root can change a file's ownership.)

Because we want the script to be maintained by minx, it would seem sensible to keep it somewhere beneath minx's home dir /home/minx – a subdirectory called cgi-bin, for example:

```
cd /home/minx/cgi-bin
```

```
ls -l dater
```

```
-rwxr-x--- 1 minx webuser 81 May 16 08:29 dater
```

To check the script works, we invoke it directly from the prompt:

```
./dater
```

```
Content-type: text/html
```

```
Content-type: text/html
```

```
Date and time are : Thu May 16 09:10:33 BST 2002
```

What we are seeing here is the output which would be returned to the browser if the script was invoked as a CGI script by *Apache*. There is one more important step before we can run our script as a CGI script. We need to add a line to httpd.conf:

```
ScriptAlias /minx-cgi/ /home/minx/cgi-bin/
```

This directive does two things. First, it defines an alias, so that requests from the browser for files inside the directory 'minx-cgi' (normally interpreted relative to the **DocumentRoot**) are interpreted relative to the directory /home/minx/cgi-bin. Second, it tells *Apache* that files within this directory are to be treated as CGI scripts – instead of simply returning the contents of the file back to the browser, *Apache* will execute the file, collect its standard output, and send that back to the browser. To access the script from the browser we simply type in its URL, like this:

```
http://localhost/minx-cgi/dater
```

and if all is well, the output from the **date** command will appear in the browser window. If you then reload the page, you should see the time update, and you have just dipped your toes into the vast ocean of dynamic web content!



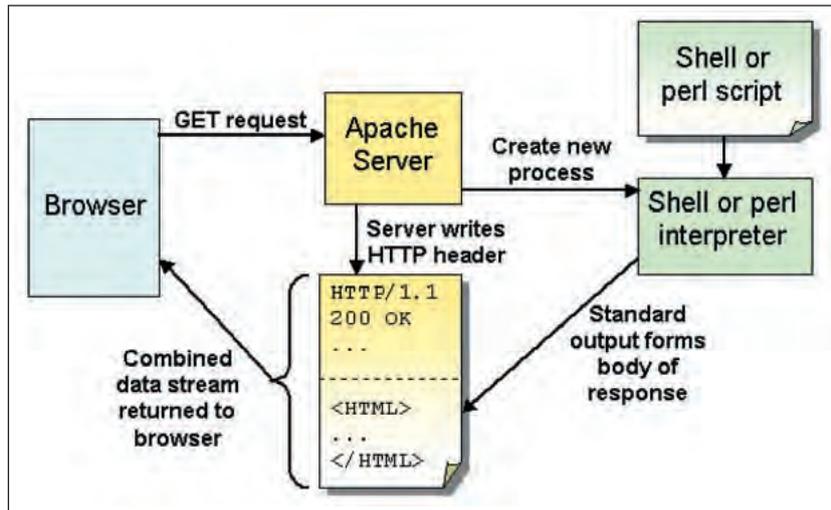


Fig 4. CGI Architecture.

## Processing Forms with CGI Scripts

In reality it's not common for a CGI program to be invoked directly by entering its URL. More often, it's invoked to process data entered by a user who fills in and submits a form on a web page. **fig. 5** shows the overall sequence of events. The form will be defined by HTML code which will look something like this:

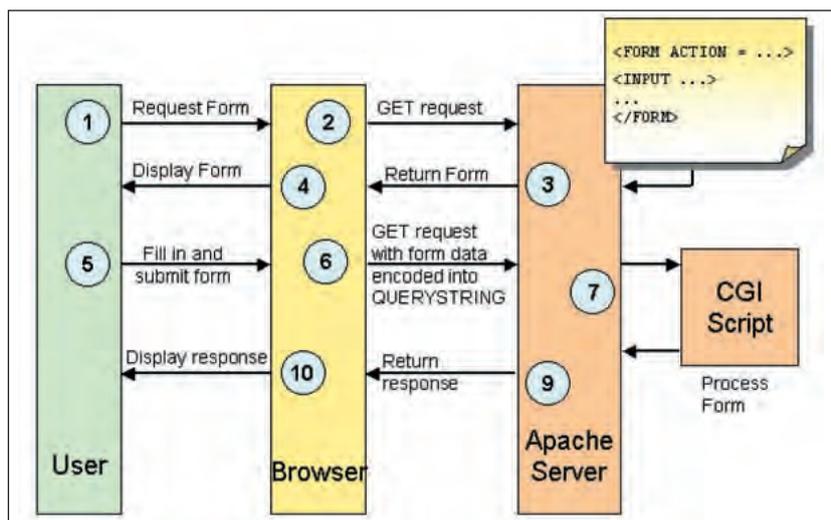
```
<FORM ACTION="/minx-cgi/grabform" METHOD="GET">
Part Number: <INPUT TYPE="text" NAME="partnum">
Quantity: <INPUT TYPE="text" NAME="qty">
<INPUT TYPE="submit" VALUE="Place Order">
</FORM>
```

When the browser sees this `<FORM>` tag, it displays a simple form with text entry boxes labelled 'Part Number' and 'Quantity' and a button labelled 'Place Order'. If the user enters text into the boxes and presses the button, the browser will issue an HTTP **GET** request for the file specified in the **ACTION** parameter within the `<FORM>` tag – in this case, the file `/minx/grabform`. This is the CGI script responsible for processing the form data. The **GET** request sent to *Apache* includes an encoded version of the data which the user entered into the text boxes, like this:

```
GET /minx/grabform?partnum=1443&qty=100 HTTP/1.1
```

*Apache* will take the encoded data (the string `partnum=1443&qty=100`) and pass it to the CGI script through an environment variable called **QUERY\_STRING**. It's left up to the script to parse the individual components and do something with the data. This

Fig 5. Sequence of events processing a form.



gets a little messy in a shell script but is easy in Perl. By the way, there is nothing to stop you manually entering a URL with an encoded query string, as shown above, to test a script written to process a form without having written the HTML for the form!

## The Bad News about CGI

If you're going to allow people to run CGI scripts on your site you need to be aware of the security implications. Remember, the script will be run with the user and group identity of the web server. If you're writing your own scripts, you presumably trust yourself to make sure they don't do anything bad. But are you willing to trust the rest of your user community? Many web admins, especially ISPs, don't allow you to run CGI scripts at all.

One of the major disadvantages of CGI is its relatively poor performance. Every time a CGI program is executed, *Apache* has to start a new process to do so. For a small, fully compiled CGI program in C, it's not too bad, but starting up the Perl interpreter (or the shell) takes considerably longer. There is a module available for *Apache 1.3* called **mod\_fastcgi** which allows CGI scripts to run using a slightly different model. In particular, instead of exiting completely and having to be restarted for each request, scripts written for **fastcgi** operate in a loop; at the top of the loop they issue a call to a library routine (provided as part of **fastcgi**) which blocks the process until the next request is received. In this way the per-request startup overhead is avoided, but the clean isolation between the script and the web server is maintained. For details, see [www.fastcgi.com](http://www.fastcgi.com)

## Introducing PHP

PHP originally stood for 'Personal Home Page', but in line with the current fashion for recursive acronyms is now short for '**PHP Hypertext Preprocessor**'. It's a scripting language designed to be integrated directly into a webpage. PHP's syntax is somewhere between C and Perl, making it easy for most Linux programmers to learn. Another of its strengths is excellent support for connecting to databases. A very popular database to use in this context is *MySQL* – the combination of Linux, *Apache*, *MySQL* and PHP makes a very potent, free collection of tools for hosting commercial web applications. But there's native support for other databases, too, including *PostgreSQL*, *mSQL*, *Oracle*, *Unix dbm*, *Informix* and *Sybase*, and any db which provides an ODBC driver.

For our PHP demonstration, we'll use the same 'print the date' example as for our CGI script. Then we'll move along and see how easy it is to process form data in PHP.

## Some Self-Assembly Required

PHP is not part of the standard *Apache* distribution, and there's a bit of work to do to get it installed and set up. It's normally built as a dynamically shared object (DSO) which is linked into *Apache* at run time. (In terms of size, PHP completely dwarfs *Apache* – on my system the *Apache* executable is 330KB and the PHP module is 2.87MB!) Building it from source is straightforward; I built mine from a source tarball I found on a coverdisc. The sequence of operations went something like this:

```
# cd /usr/local
# tar xzvf /mnt/cdrom/src/php-4.0.4.tar.gz
# cd php-4.0.4
# ./configure --with-apxs
# make
# make install
```

The `--with-apxs` option on the configure command is

needed to ensure PHP is built as a DSO. On the *Apache* side of things, there are a couple of lines to be added to `httpd.conf`:

```
LoadModule php4_module modules/libphp4.so
AddModule mod_php4.c
AddType application/x-httpd-php .php .php4
```

The first line instructs Apache to dynamically load the PHP shared object, `libphp4.so`. (The path name to this file may be different, depending on your installation.) The second line adds the module to *Apache's* active module list, and the third line tells *Apache* to associate the 'MIME' type `application/x-httpd-php` with files bearing the extensions `.php` or `.php4`. In effect this line tells *Apache* that files with these extensions should be submitted to the PHP module for parsing.

With these chores over, we can try our first PHP program:

```
<HTML>
<HEAD> <TITLE> PHP Example </TITLE> </HEAD>
<BODY>

The date is: <B>
<?php echo date("jS F Y"); ?>
</B>

</BODY>
</HTML>
```

We'll call it `dater.php`, and put it in the **DocumentRoot** directory. Now we can test it by pointing the browser at `http://localhost/dater.php`; the result is shown in **fig. 6**.

Notice how in this example, the PHP code is embedded within an HTML page by enclosing it in `<?php ... ?>` tags. Some of the output comes from the static HTML code, which can be thought of as a kind of template, and some results from executing the PHP code. The output from PHP's `date()` call is displayed in bold because of the `<B> ... </B>` tags in the surrounding HTML. This ability to mix-and-match static (HTML) and dynamic (PHP) content is very useful – in the CGI model the script is responsible for generating the entire content returned to the browser. By the way, if you ask your browser to view the source for this example, you will not see the PHP code, you'll see the text that the PHP code generated. This nicely reinforces the notion of server-side scripting – the PHP code is interpreted on the server, not in the browser.

## Accessing Form data with PHP

We'll wrap up with a demo of how to process an HTML form in PHP. We'll use a slightly modified version of the form we used in the CGI example. (I added some `<BR>` and `<P>` tags to lay the form out better.) Here's the HTML code that defines the form:

```
<HTML>
<FORM ACTION="/processform.php" METHOD="GET">
Part Number: <BR> <INPUT TYPE="text"
NAME="partnum"> <P>
Quantity: <BR> <INPUT TYPE="text" NAME="qty"> <P>
<INPUT TYPE="submit" VALUE="Place Order">
</FORM>
</HTML>
```

**Fig. 7** shows how the form appears when displayed in the browser. If the user fills in the form and presses the 'Place Order' button, the browser will, as we saw earlier, submit a **GET** command to the server with the form data encoded into it:

```
GET processform.php?partnum=1234&qty=100
HTTP/1.1
```

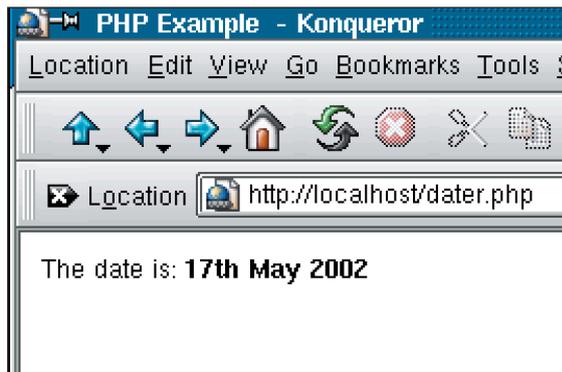


Fig 6. PHP output viewed in browser.

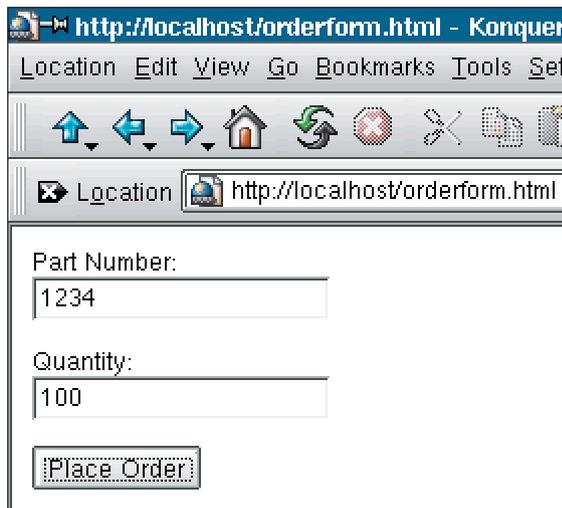


Fig 7. Order form viewed in browser.

Here's a simple PHP script to process the data.

```
<HTML>
<HEAD> <TITLE> Form Processing </TITLE> </HEAD>
<BODY>

<?php
echo "Thanks for ordering ".$qty." of part ".$partnum;
mail ("chris@localhost", "order received",
"received order for ".$qty." of part ".$partnum);
?>

</BODY>
</HTML>
```

Notice how easy it is to retrieve the values entered into the named fields on the form – you simply reference the variables `$qty` and `$partnum`, which are named after the fields specified within the `<FORM>` definition. We've also included a call to `mail()` to send a message to someone notifying them that an order has been placed, just to show how easy it is. By the way the dots you see in the above code are PHP's string-joining operator.

## Coda

Over the last four months, starting from an uninstalled tarball, we've built and configured an *Apache* server which can support IP-based and name-based virtual hosting, controlled access to web resources based on the browser's IP address, created password-protected areas of our site, logged site accesses and errors, created directory listings (both plain and fancy), and generated dynamic web content using two different languages. Now try a few experiments of your own. Good luck! [LXF](#)

## JUKEBOX PLAYERS

# mp3 madness



Too lazy to put a new disc in the CD player every hour, **Charlie Stross** turns to Perl to organise an MP3 jukebox, and clears his study of precariously stacked jewel cases.

Linux can drive a number of PC sound output devices (and sound chips on non-PC hardware, such as the AWACS sound chip on my Apple iBook). In general, you need to load a kernel module that drives your sound hardware. There are two families of sound driver; the older OSS drivers, and the new ALSA drivers. ALSA is thread-safe, multi-processor safe, and provides a legacy API so that OSS applications can use it; if given a choice, choose ALSA. Your applications usually send data to, and read data from, a sound card via a couple of special device files: /dev/dsp (a digital sampling device – you read incoming audio data from this), /dev/audio (a Sun compatible audio output device that takes .au files and makes sounds come out of your speakers), /dev/mixer (controls volume settings for each device, left/right balance, and other parameters), /dev/sequencer (if present, lets you play MIDI, FM, and other sequencer files), and /dev/sndstat or /proc/sound (which you can read sound driver settings from).

You can copy a sound stream from /dev/dsp to a file, or copy a .au file to /dev/audio and make noises: all else boils down to file format translation. (For a more detailed overview see the Sound-HOWTO: <http://tldp.org/HOWTO/Sound-HOWTO/index.html> and MP-Playing-HOWTO: <http://tldp.org/HOWTO/MP3-HOWTO.html>.)

MP3 (MPEG3) is a standard for file compression oriented towards frame based video and audio files. The Compact Disk Digital Audio (CDDA) format used on CDs is basically a raw, uncompressed bit-stream with some additional error-correction encoding, recorded in 12-bit chunks at a sampling speed of 44100 samples/second. MPEG3 applies some fairly efficient compression to the bitstream, and reduces the volume of data from the roughly 10MB/minute of a CD to roughly 1MB/minute (at high quality), or as little as 200KB/minute (AM radio broadcast quality). MP3 files are optimised for decompression, because their usual use is playback; it takes a lot longer to squish the data down than to unpack it again. The MP3 file format isn't just a straightforward stream of sound; it needs to contain two or more channels (for stereo), volume and timing info, and ID3 tags that provide meta-information about the track's creator and name.

Turning a CD into a set of MP3 tracks takes a tool such as *Grip* (from [www.nostatic.org/grip](http://www.nostatic.org/grip)), a front-end for a bunch of other tools. First in the chain is either *cdda2wav* or the more sophisticated *cdparanoia*. These tools read an audio CD and 'rip' the tracks on it into WAV files – straight digital sound streams on disk. *Lame*, *Bladeenc*, *Benc*, or any other MP3 ripper then reads a WAV file and writes a corresponding compressed MP3 file. Because CD's don't contain track and artist info in human-readable form, Grip also generates a disk-ID number for the CD and polls the Free-CDDDB server to see if anyone's entered the artist and track details for that disk. (Free-CDDDB is a centralised database, intended to save the effort of typing in all the details by

hand every time; see [www.freedb.org/modules.php?name=Sections&sop=viewarticle&artid=6](http://www.freedb.org/modules.php?name=Sections&sop=viewarticle&artid=6) for details of how it all works.

On Linux, you've got a huge choice of mp3 players – from the simple-looking command line tools like *mpg123* to hairily complex graphical players like *Xmms* (formerly *X11Amp*, a *WinAmp* clone) which can also handle streaming mp3s, ogg files, special effect filters, and act as a graphic equaliser.

I decided to do things a bit differently. Archos (see [www.archos.com](http://www.archos.com)) make a very nice toy called the Archos Jukebox Recorder 20. This is basically a portable 20GB hard disk with an mp3 decoder in firmware. You can plug it into your Linux system via a USB cable, and mount it as a USB mass storage device (which makes it look like a slow SCSI disk to the kernel). Once mounted you can copy MP3 files and M3U playlists onto it, then unplug it and leave the job of playing music up to the dedicated jukebox instead of having it bog down your computer.

## Finding files and building playlists

One of my first headaches with the jukebox was this: you can tell it to play all the tracks in a directory, but it won't dive into subdirectories. With 400 odd albums loaded onto it (thanks to a marathon session with *Grip*) I settled on the obvious file structure of having a directory for each artist or band, and a subdirectory for each album. How do you get a jukebox to play every track by an artist, if it only plays the mp3s it finds in one directory?

It turns out that *WinAmp* introduced a special type of file called a playlist. A playlist is, at its simplest, a list of filenames presented one per line; alternatively URIs may be used for remote streaming files, but if we're doing this on a jukebox all we need is relative pathnames to each file in the list. At its simplest, we can build a playlist without Perl, using a simple shell script:

```
#!/bin/sh
JUKE=/mnt/jukebox
$base=$(pwd)
for foo in $( find $JUKE -maxdepth 1 -type d -print)
do
  echo processing $foo
  NAME=$(basename $foo)
  cd $foo
  find . -type f -name "*.mp3" -print > $foo/$NAME.m3u
  cd $base
done
```

But what happens if we want a hybrid playlist containing, say, all the tracks recorded by Bjork – and The Sugarcubes (her original band)? Or all the tracks from albums she was involved in between 1989 and 1995?

Here's a first cut at the sort of thing we might do:

```
#!/usr/bin/perl
```

```

use File::Find;

my $startdir = '/home/mp3';
my $target = qr(bjork|sugarcubes);

find(\&wanted, $startdir);

sub wanted {
    if ( /\.mp3$/ && lc($File::Find::name) =~ $target ) {
        print "$File::Find::name\n";
    }
}

```

This is a fairly simple application of a very handy Perl module – **File::Find**. **File::Find** is part of the standard Perl distribution. Perl, in addition to being a sufficiently strict superset of the *sed* and *awk* text editing tools to have actual *sed*-to-Perl and *awk*-to-Perl translators, is also a superset of *find*.

*find* (used in the first example) walks a directory tree and applies a series of tests (specified on the command line) to whatever files it finds in a directory. It discards files as they fail the tests, and then applies any specified actions to the remaining files – e.g. the **-print** action prints the file's name, and the **-exec** action runs another program (you use two empty {}, to denote the filename as you interpolate them into your **-exec** command).

Perl's unary file test operators are applied to a scalar variable:

```
-d '/home/mp3'
```

Returns true if **/home/mp3** is a directory, and false otherwise. **-r**, **-w**, and **-x** test for the **read**, **write**, and **execute** attributes (in the context of the current processes' effective user ID – that is, they test whether the program can read, write or execute the file). **-z** tests if the file is of zero length, while **-s** returns the size of the file in bytes; **-T** tests whether it's a text file, and **-B** tests whether it's a binary file. You can daisy-chain file tests together:

```
if ( -r $myfile && -s $myfile && -T $myfile ) {
# $myfile is readable and of non-zero length and contains text

```

Every time you use a file test operator, Perl has to execute the **stat()** system call. You can cut down on this filesystem access by either calling **stat()** yourself in Perl (it returns an array of descriptive data relating to the file), or you can use a shortcut – Perl caches the test information for the last file you tested, and the symbol **'\_'** points to this copy:

```
if ( -r $myfile && -s _ && -T _ ) {
# same as above, but with less disk access

```

**File::Find** is a module that lets us write directory-traversal code in Perl. It provides a couple of utility subroutines, the most important of which is **find()**. **find** takes a variety of parameters, but the bare minimum is a code reference pointing to a test subroutine, and a directory to traverse. For example:

```
find(\&wanted, $startdir);
```

This tells **find()** to search all directories within **\$startdir** for files, and for every file, to run the subroutine **wanted()**, defined later.

Inside **wanted()**, we have access to some special variables.

**\$\_** is set to the name of the current file, **\$File::Find::name** to its complete pathname, and **\$File::Find::dir** is the current dir's name. **wanted()** implicitly calls **chdir()** to move into the current subdir before handing control to the code you've written:

```

if ( /\.mp3$/ && lc($File::Find::name) =~ $target ) {
    print "$File::Find::name\n";
}

```

If the filename (that's in **\$\_**) ends in **.mp3**, and the full

pathname of the file matches the quoted regular expression

**\$target**, print the complete pathname.

**\$target** contains a quoted regular expression, specified with **qr()** – everything inside the brackets is transformed into a regexp. This makes it easy for us to write programs that take regexps as parameters; we can turn strings into regexps and use them later. Here, we're just using it for clarity: our search pattern is **bjork|sugarcubes**, which matches *either bjork or sugarcubes*.

Handy tip: if you can think of a **find(1)** expression to locate a file, and want some Perl that does the same job, use the program *find2perl* to do the Perl for you. E.g. to write a Perl script like:

```
find /home/mp3 -type f -mtime -3 -name '*.mp3' -print
```

(This means: find all items of type 'regular file' with a modification time less than three days ago and a name matching **\*.mp3**, then print their names.) You can substitute **find2perl** for **find** in the above line and it will spit out something like this:

```

#!/usr/bin/perl -w
eval 'exec /usr/bin/perl -S $0 ${1+"$@"}'
if 0; # $running_under_some_shell

use strict;
use File::Find ();

# Set the variable $File::Find::dont_use_nlink if you're using
# AFS,
# since AFS cheats.

# for the convenience of &wanted calls, including -eval
# statements:
use vars qw($name $dir $prune);
$name = *File::Find::name;
$dir = *File::Find::dir;
$prune = *File::Find::prune;

# Traverse desired filesystems
File::Find::find({wanted => \&wanted}, '/home/mp3');
exit;

sub wanted {
    my ($dev,$ino,$mode,$nlink,$uid,$gid);

    (($dev,$ino,$mode,$nlink,$uid,$gid) = lstat($_)) &&
    -f _ &&
    (int(-M _) < 3) &&
    /\.mp3$/s &&
    print("$name\n");
}

```

This shows us how to traverse a directory tree and build a list of files – if we just modify that final print statement to print to a previously opened file handle, we can turn its output into an m3u playlist containing only those mp3 files modified in the past three days. And by using regular expressions we can pick only mp3 files with the right name. But is that the best way to do things?

## Chasing ID3 tags

Names of files are not inextricably related to their contents. I have no guarantee that my mp3 ripper correctly connected to the FreeCDB service and correctly retrieved the name of the files it was pulling off the CD; I might have a potload of directories called 'untitled artist' containing subdirectories called 'untitled album' filled with files called '01.mp3, 02.mp3 ...' and so on.



And we certainly can't expect a simple-minded **find** script to pull in all tracks from albums Bjork participated in between 1989 and 1995, because the necessary information (about when the album was recorded) simply isn't in the filename.

There is a way to name mp3s so that you can't lose the info simply by renaming the file. mp3 files contain ID tags – specifically, ID3 tags, a standard record structure embedded in the header of the file that indicates the artist, album, track name, year, copyright status, and other important information. The ID3 tags are also used by MP3 players (like the Archos jukebox) to provide captioning information during playback.

Let's look at two ways of setting or examining ID3 tags– the command line utility *id3tool* (<http://kitsumi.xware.cx/id3tool/>), and the Perl modules in **MP3::Tag** (**MP3::Tag::ID3v1** and **MP3::Tag::ID3v2**). *id3tool* is a Version 1 ID3 tagfile editor; **MP3::Tag** can handle v2 tags as well. Install them off CPAN:

```
perl -MCPAN -e 'install MP3::Tag;'
```

Let's look at *id3tool* first. To set the Version 1 ID3 tags on a file:

```
id3tool -t 'World in my Eyes' -a 'Violator' -r 'Depeche Mode'
01-world\ in\ my\ eyes.mp3
```

To read the tags from a file:

```
id3tool 01-world\ in\ my\ eyes.mp3
```

```
Filename: 01-world in my eyes.mp3
```

```
Song Title: World in my Eyes
```

```
Artist: Depeche Mode
```

```
Album: Violator
```

(Only ID3 information actually in a file gets returned by *id3tool*.)

We can use *id3tool* from within Perl to get info about mp3 files:

```
my $file = '01-world in my eyes.mp3';
```

```
my $info = `id3tool $file`;
```

```
chomp $info;
```

```
my @info = split(/\n/, $info);
```

```
my %info = ();
```

```
foreach (@info) {
```

```
my ($k, $v) = split(/:|s+/, $_);
```

```
$v =~ s/^\s+//;
```

```
$info{$k} = $v;
```

```
}
```

```
print "Title:", $info{'Song Title'}, "\n";
```

```
# and so on
```

Each time we want to look at a new file we have to spawn a subshell to run an external program. **MP3::Tags** is more efficient:

```
use MP3::Tag;
```

```
$mp3 = MP3::Tag->new($filename);
```

```
# get some information about the file in the easiest way
```

```
($song, $track, $artist, $album) = $mp3->autoinfo();
```

```
# or have a closer look on the tags
```

```
# scan file for existing tags
```

```
$mp3->get_tags;
```

```
if (exists $mp3->{ID3v1}) {
```

```
# do something with them
```

```
$id3v1 = $mp3->{ID3v1}; # a handy shortcut
```

```
print " Song: " . $id3v1->song . "\n";
```

```
print " Artist: " . $id3v1->artist . "\n";
```

```
print " Album: " . $id3v1->album . "\n";
```

```
print " Comment: " . $id3v1->comment . "\n";
```

```
print " Year: " . $id3v1->year . "\n";
print " Genre: " . $id3v1->genre . "\n";
print " Track: " . $id3v1->track . "\n";
if (!exists $mp3->{ID3v2}) {
# create a new ID3v2 tag
$mp3->new_tag("ID3v2");
# create ID3v2 album title
$mp3->{ID3v2}->add_frame("TALB", $id3v1->album);
# write out the ID3v2 tag
$mp3->write_tag;
}
}
```

```
$mp3->close();
```

Version 1 ID3 tags are fields with preset names. Version 2 added the ability to define arbitrary tag names: a *frame* is a tag and its value, and tags are identified by four-character codes. There are simple frames (name/value pairs) and complex frames that may contain hashes of additional data. Get a list of simple frames supported by **MP3::Tag** by calling **\$mp3->id3v2->get\_frame()**.

ID3v2 is extensible, and has a whole load of handy frames available. *E.g.* a frame for an involved people list (IPLS); if we've got a properly tagged album there are tags for Lead performer/ Soloist (TPE1), Conductor/performer (TPE3), Lyricist (TEXT), Composer (TCOM), *et al.* There are also tags that tell us how the mp3 is being used – Play counter (PCNT), and TPUB (publisher) or WPAY (payment). On an mp3 server, see how popular our tracks are by splicing the following code into our scripts:

```
use MP3::Tag;
```

```
my $mp3 = MP3::Tag->new($filename);
```

```
$mp3->get_tags;
```

```
if (!exists $mp3->{ID3v2}) {
```

```
$id3v2 = $mp3->new_tag("ID3v2");
```

```
$id3v2->add_frame("PCNT", 1);
```

```
} else {
```

```
my ($info, $name) = $id3v2->get_frame("PCNT");
```

```
$info++;
```

```
$change_frame("PCNT", $info);
```

```
}
```

```
$id3v2->write(tag);
```

(This opens the mp3 file **\$filename** and gets the tags. If no v2 tags exist, it creates them, adds 1 to the frame PCNT, and saves it. If they exist, it retrieves the value of PCNT, increments it, and re-saves it. Each time we call this code, the counter is incremented.)

*Grip* doesn't automatically write ID3 tags when it rips files. (Nor does Apple's *iTunes 2.0*, for that matter.) How can we easily build ID3 tags for a directory tree full of files we just ripped?

It turns out that **MP3::Tag** is smart enough that it can try and yank the info out of the file's pathname if no tags can be found. To configure it to do this, we call **MP3::Tag->config()**, and tell it to use a specific search order among ID3v1 tags, ID3v2 tags, or the filename, when returning information with **autoinfo()**. We then use **autoinfo** to get our basic info, and create some tags:

```
#!/usr/bin/perl
```

```
use MP3::Tag;
```

```
my $file = shift @ARGV;
```

```

MP3::Tag->config("autainfo", "filename", "ID3v1", "ID3v2");

my $mp3 = MP3::Tag->new($file);

my ($song, $track, $artist, $album) = $mp3->autainfo();

print "song: $song\ntrack: $track\nartist: $artist\nalbum:
$album\n\n";

if (! exists $mp3->{ID3v2} ) {
    $mp3->new_tag("ID3v2");
    my $id3v2 = $mp3->{ID3v2};
    $id3v2->add_frame("TALB", $album);
    $id3v2->add_frame("TOPE", $artist);
    $id3v2->add_frame("TRCK", $track);
    $id3v2->add_frame("TIT2", $song);
    $id3v2->write_tag();
}
if (! exists $mp3->{ID3v1} ) {
    $mp3->new_tag("ID3v1");
    my $id3v1 = $mp3->{ID3v1};
    $id3v1->song($song);
    $id3v1->artist($artist);
    $id3v1->album($album);
    $id3v1->track($track);
    $id3v1->write_tag();
}
$mp3->close;

```

This won't get you the interesting year info, but it helps with the basics by building you a rudimentary set of ID3 tags even if your ripper doesn't do it for you. For that you need to go to CDDB.

## Querying CDDB in Perl

The FreeDB and CDDB databases provide a network interface to allow client programs to talk to them; and if you're using Perl, a variety of modules are available. First, there's *CDDB.pm* itself. This module is a formal client-side implementation of the CDDB API. Given an actual music CD's table-of-contents data, it composes a query and either retrieves the CDDB records for corresponding disks, or provides facilities for submitting a CDDB entry.

This is not a lot of use if what you have is a collection of mp3s with partial ID3 tags, and what you want to do is flesh out the information held on them. For example, let's go back to Bjork. Using **MP3::Tag**, we've created ID3 tags for all our MP3's – but they don't include release year or some of the supplementary information we want. To solve the problem we need a couple of tools – one to find the CDDB entry number for an album, and one to retrieve the CDDB entry and parse it.

You can get an album's CDDB ID via the web-based search interface to CDDB servers. For example, you can search FreeDB via the URL [http://www.freedb.org/freedb\\_search.php](http://www.freedb.org/freedb_search.php), which takes parameters such as **words**; a **GET** request for:

```
http://www.freedb.org/freedb\_search.php?words=bjork
```

Retrieves an HTML page containing links to a URL like:

```
http://www.freedb.org/freedb\_search\_fmt.php?cat=blues&id=bf11350e/
```

With a text caption of the artist and album title.

This makes it possible for us to yank a CDDB record for a specific named album, but as with much code for searching online databases it gets messy fast. First, we use **LWP::Simple** to fire a web request at [www.freedb.org/freedb\\_search.php](http://www.freedb.org/freedb_search.php); then we

run the response through **HTML::LinkExtr** to extract the HREF links. We look through these for links that look like hits – ones with **cat=(category)&id=(number)** appended to them. We then use **LWP::Simple** to retrieve the appropriate FreeDB entry file from the server, **CDDB::File** can parse it and turn it into a Perl object that we can work with. For an encore, try adding code to validate these search results – the database may give more hits than you expect for a given artist and album title. Once you've done it, typically by counting track numbers and durations, the way is clear to build ID3 tags using them. [LXF](#)

```

#!/usr/bin/perl

use LWP::Simple;
use HTML::LinkExtr;
use CDDB::File;

my $artist = 'depeche mode';
my $record = 'Master and servant'; # "personal jesus";
my $cddb = "http://www.freedb.org/freedb_search.php";
my $words = $artist . " " . $record;
my $response = s/s/+g;
my $request = "";

my $doc =
get("$cddb?allfields=YES&allcats=YES&words=$words");

# now we need to parse the result file for links and extract
CDDB index numbers

my $p = HTML::LinkExtr->new();
$p->parse($doc);
my @links = $p->links;
@links = grep { $_->[2] =~ /\&id=/ } @links; # get rid of
non-relevant URLs
@links = map { $_ = $_->[2] } @links; # get rid of
spurious fields
# @links is now an array of URIs into CDDB. Let's turn 'em
into entry ID numbers

foreach (@links) {
    $cat = $id = $_;
    $cat =~ s/^.cat=(.*)\&.*$/1/;
    $id =~ s/^.id=(\d{abcdf}+)$/1/;
    # now we have a category and an ID we can retrieve the
CDDB data file
    my $request = "http://www.freedb.org/freedb/$cat/$id";
    $doc = scalar get($request);
    open (OUT, ">$cat.$id.tmp");
    print OUT $doc;
    close OUT;
    my $cddb = CDDB::File->new("$cat.$id.tmp");
    unlink "$cat.$id.tmp";
    print "Category: ", $cat, "\n";
    print "CDDB ID: ", $id, "\n";
    print "Artist is: ", $cddb->artist, "\n";
    print "Disk is: ", $cddb->title, "\n";
    print "Year was: ", $cddb->year, "\n";
}
exit;

```

## BUILDING OBJECTS FROM XML

# Speaking Java



Parsing XML is only part of the story. How do you make use of that information from within Java? **Richard Drummond** investigates.

**W**here are we at? Oh, yes. We began looking at how you can parse XML data in Java with SAX, the simple API for XML. Last time we introduced the basics of SAX with a view to marshaling data stored as XML, in this case, an XML file containing an index of the files stored on a *Linux Format* coverdisc. You should now know how to use SAX to extract the data stored in an XML file; the next question is how do we make use of that data? How do we convert XML into Java objects?

If you remember, SAX is a very simple API which simply processes an XML stream sequentially, allowing us to pluck out the values we are interested in by implementing a system of callbacks. The other principal XML API for Java – DOM or the Document Object Model – actually converts the hierarchy of data stored in an XML file into a tree of Java objects. This is easy to use, but has a high price in performance and memory overheads. Hence we are using SAX and have to do things manually.

## Building models

Before we can instantiate objects, we need to think about what entities those objects will be encapsulating. If we consider the types of data we will be dealing with in our example project, the *LXF* CD Index browser, we can classify the data into two types of object: packages and files. In our model, a package is a particular piece of software stored on a specific disk. Thus each package object has with it an associated name, for example *Apache* or *Mozilla*; a version number; a description; the name of the disc that it is located on, for example, *LXFCD23*; a URL of the project home page; and the path to the directory where it is stored on that disc. We quite often provide the same piece of software in different distribution formats such as a source tarball, binary tarball, RPM or Debian package. Thus a 'file' is the file containing a particular distribution format of some 'package' and a single 'package' will have one or more files associated with it. Each file has a name, that is, its filename relative to the path of the project it belongs to; a MIME type; and a size. We could model these two types of object in Java with the following classes:

```
class IndexedPackage
{
    String Name; /* The package name */
    String Disc; /* Which disc was it on? */
    String Description; /* A brief description */
    String Path; /* Path to directory on disc */
    String Version; /* Version number */
    String HomePage; /* Website URL */
    ArrayList Files; /* This package's files */

    // the rest of the class goes here . . .
}
```

```
class IndexedFile
{
    String Name; /* The file name */
    String Type; /* It's MIME type */
    String size; /* It's size */

    // the rest of the class goes here . . .
}
```

I have not been too particular about the how attributes associated with each class have been represented internally as fields; for now, each is modelled by a Java **String**. This is a good fit for most of the data we are dealing with because it is mostly text data, but it might actually be more useful or more efficient, for instance, to store a file's size as an **int** or a **long**, or to store a package's home page as a **URL**. Since the class mechanism hides this information from the user of the class, we are free to change these representations at a later date if we so choose.

Armed with these classes, we can tackle the job of building objects using the data contained in an XML file. Actually, for the rest of the article, to keep things simple and to save space, we will only consider constructing **IndexedPackage** objects. We can build a list of **IndexedFile** objects similarly, but we'll leave that until next time.

## Maintaining state

Some of the data we are interested in is stored as element attributes according to our XML schema, while the rest is stored as element data. The **startElement()** method of SAX's **ContentHandler** interface provides access to the former, while the **characters()** method provides access to the latter. The **characters()** method, if you recall, simply hands us chunks of character data in an array as it is parsed – we are given no clue as to the context of that data. Thus to be able to make use of that character data, we need to be able to keep track of the state of the parser – that is, what XML element is currently being parsed – ourselves. When we know the state, we know what element a piece of data belongs to.

It's a fairly simple task to track the state of the parser. We can do this with a member variable whose current value corresponds to the current XML element. (We set up some mapping between XML elements and integers, say.) When we encounter an open or close tag – signalled by a call to our **startElement()** and **endElement()** methods, respectively – we bump the value stored in the state variable accordingly.

Here is the modified **startElement()** method of our **IndexParser** class:

```
public void startElement( String uri, String name, String
qName, Attributes attributes ) throws SAXException
{
    switch ( state ) {
```

```

case STATE_START:
if( name.equals( "indexedfiles" ) )
state = STATE_INDEX;
else
throw new SAXException( "Parse error" );
break;
case STATE_INDEX:
if( name.equals( "package" ) )
{
state = STATE_PACKAGE;
packageName = attributes.getValue( 0 );
}
else
throw new SAXException( "Parse error" );
break;
case STATE_PACKAGE:
if( name.equals( "disk" ) ) state = STATE_DISK;
else if( name.equals( "path" ) ) state = STATE_PATH;
else if( name.equals( "link" ) ) state = STATE_LINK;
else if( name.equals( "version" ) ) state = STATE_VERSION;
else if( name.equals( "description" ) ) state =
STATE_DESCRIPTION;
break;
}
}

```

If you examine the above code, you can see an immediate and added benefit of tracking the parser's state: we can do a rudimentary amount of semantic checking ourselves. When SAX is run with validation turned off, it is faster, but the parser no longer cares whether the XML document being parsed matches the grammar specified by the DTD (Document Type Definition). However, if we track state, we know that only certain state transitions are permissible, and so can flag errors when an unexpected XML element occurs.

As well as maintaining state when opening tags are encountered, the above code extracts the values of any element attributes that we are interested in. This how we find a package's name: it's the first attribute of the `<package>` element. We store this in an instance variable for later use.

Next, our `characters()` method now becomes:

```

public void characters( char[] ch, int start, int length ) throws
SAXException
{
String tmpstr = new String( ch, start, length );

switch ( state ) {
case STATE_DISK:
diskName = diskName + tmpstr; break;
case STATE_PATH:
packagePath = packagePath + tmpstr; break;
case STATE_LINK:
packageLink = packageLink + tmpstr; break;
case STATE_VERSION:
packageVersion = packageVersion + tmpstr; break;
case STATE_DESCRIPTION:
packageDesc = packageDesc + tmpstr; break;
}
}

```

Here we take the sub-string that the parser gives us and append it to a corresponding temporary string according to the current state. As we said last time, the parser can pass data to us

in multiple chunks, so we only know that we got the complete string associated with a particular element's data when that element's end tag is encountered. We use these temporary variables – instance variables of our class – to build these chunks into complete strings.

## Ending tags

To finish the story, here is the new `endElement()` method which tracks state transitions and complements the `startElement()` method above. This also takes care of instantiating new `IndexedPackage` objects with the data we have extracted.

```

public void endElement( String uri, String name, String
qName ) throws SAXException
{
switch ( state ) {
case STATE_INDEX:
if( name.equals( "indexedfiles" ) ) state = STATE_START;
break;
case STATE_PACKAGE:
if( name.equals( "package" ) ) {
state = STATE_INDEX;
// create Package with current parsed values
IndexedPackage pkg = new IndexedPackage();
pkg.setName( packageName );
pkg.setVersion( packageVersion );
pkg.setDisc( diskName );
pkg.setDescription( packageDesc );
pkg.setPath( packagePath );
pkg.setLink( packageLink );
// add it to our list
packageList.add( pkg );
// reset temporary variables
packageName = "";
diskName = "";
packagePath = "";
packageLink = "";
packageVersion = "";
packageDesc = "";
}
break;
case STATE_DISK:
if( name.equals( "disk" ) ) state = STATE_PACKAGE; break;
case STATE_PATH:
if( name.equals( "path" ) ) state = STATE_PACKAGE; break;
case STATE_LINK:
if( name.equals( "link" ) ) state = STATE_PACKAGE; break;
case STATE_VERSION:
if( name.equals( "version" ) ) state = STATE_PACKAGE;
break;
case STATE_DESCRIPTION:
if( name.equals( "description" ) ) state = STATE_PACKAGE;
break;
}
}

```

When we encounter the `</package>` tag, we know that the parser should have retrieved enough information for us to be able to instantiate an `IndexedPackage` object to encapsulate that package's data. So, we create the object and use its accessor methods to fill in the current values. We then add it to the current list of packages, and then reset the temporary variables ready to parse the next `<package>` element. [LXF](#)

## NEXT MONTH

We didn't have time to get the source code accompanying this article on this issue's coverdisc, so we will be including it next time around. The complete code for this month uses the above methods to parse an index and build a collection of `IndexedPackage` objects from it in memory. We will add to that the rest of the code to build and handle a list of files for each package, as well as taking a look at pattern-matching with regular expressions. After all, when performing a query on a series of index files, we don't need to instantiate an object for every single package in the index – we just need to load those that match the query.



MORE GRAPHICS

# Animation effects

**PART 11:** This month **Brian Long** gets things moving with some simple animation.



Last month we saw some simple graphics using the **TCanvas** class. This month we are moving on a little to see how to achieve some simple animation effects in *Kylx* apps.

As was the case last month, the goal here is not to make a start on a professional graphics/animation package but more to introduce you to some techniques you can use to achieve some pleasing effects in your *Kylx* applications. Hopefully, *Kylx* developers can use these ideas as starting points that will lead to more interesting programs. The goal this month is to get something akin to a sprite that we can move around the form.

## The plan

Let's get some basics out of the way first. As with last time, we must use a background bitmap to build up the scene that we want on the form, rather than drawing it directly onto the form's canvas. This process is usually called "double buffering" and is commonplace in most forms of graphics application.

In fact more than that, we will probably need two background bitmaps, bearing in mind we need to move a sprite around. One bitmap can hold the intact, original rendering of the background while the other will be used to build up the scene to go on the form. Each time we need to move the sprite, we can start by copying the intact background onto the scene bitmap and then draw the sprite on top of it.

## Drawing a non-rectangular image

The first challenge we should consider is how to copy a non-rectangular image onto a bitmap. The normal canvas methods available seem to work on the basis of a rectangular region to copy from or to. Of course this is because a rectangle is a simple shape to describe, but when dealing with sprites the images are rarely rectangular.

We need to know how to treat various pixels in a rectangular region as transparent so that when we copy the image, only the relevant pixels make it to the destination bitmap and the background pixels are "lost" *en route*.

This sort of technique is commonplace in animation and can be achieved using a mask. We'll use an example of drawing an image of a sun on a form that has a coloured image drawn

across its background to show the idea. The simple project *SunDemo.dpr* consists of a form with an image component stretched across its background displaying a colourful bitmap.

When you click on the form background a sun will be drawn with its top left at that point.

The general plan involves using two bitmaps for each sprite, the image and its mask. The mask has a white background and each pixel of the actual picture is coloured black (**figure 1**). It is drawn on the destination using the **AND** operator (the value of the destination and mask pixels are combined with a binary **AND**), resulting in a black representation of the target image being left on the destination.

The demo form contains a non-visible image containing the sun image mask. When the user clicks on the form (well, actually the image on the form) this code sets the appropriate copy mode (how each pixel is copied/merged from the source to the destination) and draws the bitmap represented by the image component onto the form's canvas. The result is shown in **fig 2**.

```
Canvas.CopyMode := cmSrcAnd;
```

```
Canvas.Draw(X, Y, imgSunMask.Picture.Bitmap);
```

The sun image itself has a black background (**figure 3**) and is drawn on the destination with the **OR** operator. This lets the colour pixels go into the black masked area. An additional image component is used to hold the sun image drawn as follows:

```
Canvas.CopyMode := cmSrcPaint;
```

```
Canvas.Draw(X, Y, imgSun.Picture.Bitmap);
```

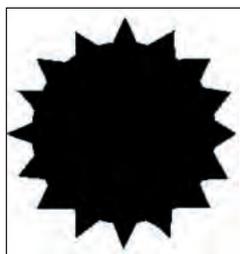


Figure 1: **Mask for an image of the sun.**



Figure 2: **The mask being used.**

The result of these two steps is just what is needed – the image is drawn, but its surrounding background obliterates any of the image on the form (see **figure 4**).

Whilst this process is reasonably straightforward, it can be tedious when you have many sprites to deal with. Things can be simplified somewhat by using an image list component (from the Common Controls page of the **Component Palette**). You can set the **Width** and **Height** properties of this to the dimensions of your sprites and then add them at design-time (using a component editor you get by double-clicking the component) or in code. The design-time editor lets you add pictures and automatically works out which colour should be considered transparent by examining the bottom left pixel (see **figure 5**).

The image list has a **Masked** property that defaults to **True** indicating that a mask will automatically be generated and used when drawing the images. The same non-rectangular image can now be drawn with this simpler single statement (used in the project *SunDemo2.dpr*) instead of the previous four statements:

```
ImageList1.Draw(Canvas, X, Y, 0)
```

## Simple animation

Now we can press on and get things moving (literally). Our first animation demo, using the tiled Union Jack background from last month, can be found in *SunAnimation.dpr*. When the app starts running the sun image can be seen in the top left of the form.

This is the form's **OnCreate** event handler.

```
type
  TfrmMain = class(TForm)
  ...
private
  BkGnd, IntactBkGnd: TBitmap; //The background and
  background backup bitmaps
  SunPos: TPoint; //defaults to (0, 0)
  procedure SetupBackground;
  function SunRect(TopLeft: TPoint): TRect;
public
  function MoveSun(X, Y: Integer): TRect;
  function DrawSun: TRect;
end;
...
procedure TfrmMain.FormCreate(Sender: TObject);
begin
  BkGnd := TBitmap.Create;
  IntactBkGnd := TBitmap.Create;
  SetupBackground; //draw offscreen background (&
  background backup)
  DrawSun; //start with sun at top left of form
  Timer1.Enabled := True;
end;
```

**BkGnd** is a bitmap used to build up each new scene to be displayed on the form and **IntactBkGnd** is the intact background bitmap. The **SetupBackground** method uses the code from last month to produce the tiled background on both **BkGnd** and **IntactBkGnd**. The next step is to draw the sun in the default position of **(0, 0)**, the top left of the form, then enable a timer.

Generally the **MoveSun** method is used to alter the record of the sun's position before calling **DrawSun**, but when the program starts the call to **DrawSun** uses the start position.

```
function TfrmMain.DrawSun: TRect;
begin
  //Draw sun in new position
```

```
  Imgs.Draw(BkGnd.Canvas, SunPos.X, SunPos.Y, 0);
  Result := SunRect(SunPos);
end;

function TfrmMain.SunRect(TopLeft: TPoint): TRect;
begin
  //
  Result := Rect(TopLeft.X, TopLeft.Y,
  TopLeft.X + Imgs.Width, TopLeft.Y + Imgs.Height);
end;
```

**DrawSun** draws the sun on the background and returns the rectangular area occupied by it (we'll see why later). When the form finishes initialising and draws itself for the very first time the **OnPaint** event will be triggered, which copies the background bitmap onto the form:

```
procedure
  TfrmMain.FormPaint(Sender:
  TObject);
begin
  //Copy as much of the bitmap as will fit on the form
  Canvas.CopyRect(ClientRect, BkGnd.Canvas, ClientRect)
end;
```

This app does no automatic movement of the image, but instead relies on the user pressing the cursor keys to control its position. The timer that was enabled earlier by the **OnCreate** event handler has an **OnTimer** event that is used to check for the keystrokes of interest at regular intervals. See the keystrokes box for more details about checking for keystrokes in this type of app, and for the code to the **IsKeyDown** helper routine.

```
uses
  Xlib;
...
procedure TfrmMain.Timer1Timer(Sender: TObject);
var
  X, Y: Integer;
  EraseRect, DrawRect, UpdateRect: TRect;
const
  XInc = 5;
  YInc = 5;
begin
  X := 0;
  Y := 0;
  if IsKeyDown(XK_Left)
  then
    Dec(X, XInc);
  if IsKeyDown(XK_Right)
  then
    Inc(X, XInc);
  if IsKeyDown(XK_Up)
  then
    Dec(Y, YInc);
  if IsKeyDown(XK_Down)
  then
    Inc(Y, YInc);
  if (X <> 0) or (Y <> 0)
  then
    begin
```



Figure 3: The sun image.

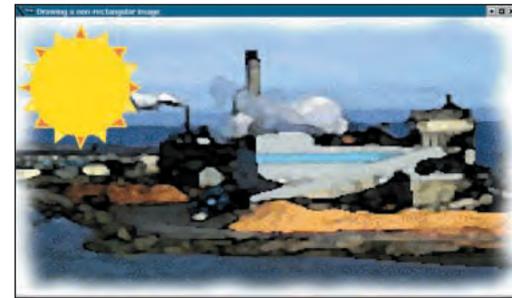


Figure 4: The final non-rectangular sun image.

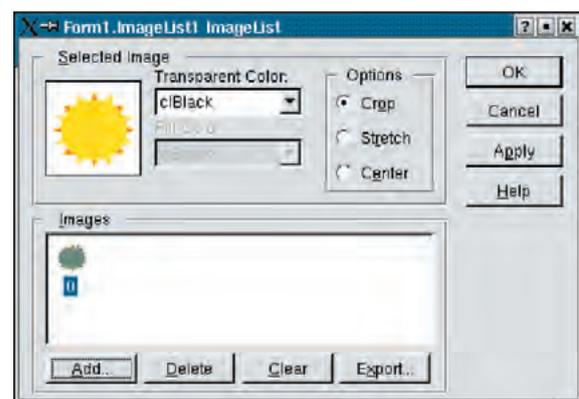


Figure 5: Adding masked images to an image list.

Figure 6: **The sun and the Union Jack – our first animation program!**



```

<< EraseRect := MoveSun(X, Y);
    DrawRect := DrawSun;
    UnionRect(UpdateRect, EraseRect, DrawRect);
    Canvas.CopyRect(UpdateRect, BkGnd.Canvas, UpdateRect);
end;
end;

```

If any of the four cursor keys are down when the timer ticks then the **X** and **Y** variables are updated. Because the checks are not exclusive this allows diagonal movement (for example if the down and right cursor keys are pressed). Once we know how the sun should be moved the **MoveSun** method is called with **X** and **Y** passed as parameters.

```

function TfrmMain.MoveSun(X, Y: Integer): TRect;
begin
    Result := SunRect(SunPos);
    Inc(SunPos.X, X);
    Inc(SunPos.Y, Y);
    //Overdraw last sun image with original background
    BkGnd.Canvas.CopyMode := cmSrcCopy;
    BkGnd.Canvas.CopyRect(Result, IntactBkGnd.Canvas, Result);
end;

```

As the sun is moved the previous image needs to be erased. The first thing **MoveSun** does is to set up the previous sun image rectangle as the return value. Next it updates the current sun position **TPoint** record and then copies a section from the intact background to cover up the old sun in the background image.

Going back to the **OnTimer** event handler, the rectangle of the old sun is recorded in the **EraseRect** variable, whilst the area covered by the new sun (returned from **DrawSun**) is then stored in the **DrawRect** variable. Since the sun only moves 5 pixels in any direction (as per the **XInc** and **YInc** constants in the timer event handler) the **EraseRect** and **DrawRect** rectangles will often overlap to a large extent. To avoid drawing a large area twice a new rectangle is calculated from the union of both these rectangles – this tells us the area on the form that needs updating to reflect the changes made by moving the sun. An appropriate chunk of the background image is then copied to the form so the user sees the sun move. You can see a snapshot of the running application in **figure 6**.

Note that our original prototypes of this code, below, didn't use this minimalist copying approach:

```

if (X <> 0) or (Y <> 0) then
begin
    MoveSun(X, Y);
    DrawSun;
    Invalidate
end;

```

This removed the old sun, drew a new one and then told the

form to redraw itself (in its entirety). The problem with doing this very regularly (one hundred times a second for responsiveness) is that it gets unacceptably “flickery”. It is always wise to try and do as little onscreen updating as you can get away with to ensure the application doesn't lose performance unnecessarily.

## Constant movement

Also on the disc is *SunAnimation2.dpr*. This version differs from the previous one in that it has no keyboard response built in. Instead the sun constantly moves around the form of its own accord, bouncing off the edges of the form when it encounters them, like a simple bouncing ball.

In fact the project source is almost identical to the previous one; the only differences are found in the **OnTimer** event handler and the **MoveSun** method.

```

procedure TfrmMain.Timer1Timer(Sender: TObject);
var
    EraseRect, DrawRect, UpdateRect: TRect;
begin
    EraseRect := MoveSun;
    DrawRect := DrawSun;
    UnionRect(UpdateRect, EraseRect, DrawRect);
    Canvas.CopyRect(UpdateRect, BkGnd.Canvas, UpdateRect);
end;

var
    XInc: Integer = 5;
    YInc: Integer = 5;

function TfrmMain.MoveSun: TRect;
begin
    Result := SunRect(SunPos);
    if ((XInc > 0) and (Result.Right > ClientWidth)) or
        ((XInc < 0) and (SunPos.X < 0)) then
        XInc := -XInc;
    Inc(SunPos.X, XInc);
    if SunPos.X > ClientWidth then
        SunPos.X := ClientWidth;
    if ((YInc > 0) and (Result.Bottom > ClientHeight)) or
        ((YInc < 0) and (SunPos.Y < 0)) then
        YInc := -YInc;
    Inc(SunPos.Y, YInc);
    if SunPos.Y > ClientHeight then
        SunPos.Y := ClientHeight;
    //Overdraw last sun image with original background
    BkGnd.Canvas.CopyMode := cmSrcCopy;
    BkGnd.Canvas.CopyRect(Result, IntactBkGnd.Canvas, Result);
end;

```

The timer has no need to check any keys – it just moves and draws the sun each time its event triggers. **MoveSun** uses a couple of variables to decide which direction to move the sun in. When one of its edges hits an edge of the form the pertinent direction is changed so it appears to bounce off.

It involves only simple code changes but gives a nice effect – we like watching the sun bouncing around the form :-)

## Summary

As you can see, moving images around a form, even with an uneven background, is reasonably straightforward. You can adapt this code and generalise it to allow control of multiple sprites, for example in a game scenario. [LXF](#)

## About Brian Long

Brian Long is a UK-based freelance trainer and problem solver for Borland's *Kylx*, *Delphi* and *C++Builder* packages. His website is at [www.blong.com](http://www.blong.com) and he can be emailed at [brian@blong.com](mailto:brian@blong.com)

## NEXT MONTH

Next month we'll look at more useful *Kylx* techniques. In the meantime, if there is something about *Kylx Open Edition* you want to see covered here, drop us an email and we'll try our best to incorporate it into a future instalment.

## Keyboard handling

Responding to keystrokes in applications and games

### APPLICATIONS

When developing a normal GUI application in *Kylux* it is typical to respond to key presses using **OnKeyDown**, **OnKeyUp** and **OnKeyPress** events. If the user presses a key on the keyboard, the **OnKeyDown** event fires for the key; if they keep the key down long enough without releasing it (half a second or so) the auto-repeat nature of the keyboard will cause additional **OnKeyDown** events to occur periodically. When the key is eventually released an **OnKeyUp** event is triggered for the key.

Note that these two events are related to physical keys on the keyboard as opposed to characters. They are triggered for every key that is understood by the X environment. Each event is passed a key code indicating which physical key was pressed or released. Constants for these are defined in the *Qt* unit and you can find them listed in the help by looking up *key codes* (fig 7).

In the case of keys that happen to generate characters the **OnKeyDown** event is immediately followed by the **OnKeyPress** event which is passed a **Char** parameter indicating the character generated.

*E.g.* if a user presses **Shift+P** (for an uppercase **P**), events triggered by the control with focus would be: **OnKeyDown** with a **Key** parameter of **Key\_Shift** as the user presses **Shift**

**OnKeyDown** with a **Key** parameter of **Key\_P** as the user presses **P**

**OnKeyPress** with a **Key** parameter of **P** to represent the generated character

**OnKeyUp** with a **Key** parameter of **Key\_P** as the user releases **P**

**OnKeyUp** with a **Key** parameter of **Key\_Shift** as the user releases **Shift**

Most controls have these events but the form can usurp them all if needed with events of its own, which are triggered in advance of those of the active control by setting the **KeyPreview** property to **True**. This allows the form to filter various keystrokes that are of form-wide interest.

Keystroke filtering is achieved by modifying the **Key** parameter in the event handler (either of the form or the control). Setting it to **0** (in **OnKeyDown/Up**) or **#0** (in **OnKeyPress**) means the active control will not see the key event. As a simple example a listbox component could ignore all up arrow presses with this:

```
procedure TForm1.ListBox1KeyDown(Sender: TObject;
var Key: Word;
Shift: TShiftState);
begin
if Key = Key_Up then
Key := 0
end;
```

You can see that there is another parameter, **Shift**, of type **TShiftState**, that is also passed to each of these events. This set indicates whether any of the **Shift**, **Ctrl** or **Alt** modifiers are being pressed. *E.g.*, if you wanted to prevent an edit control from seeing the keystrokes **Shift+X** you might write an event handler like this:

```
procedure TForm1.Edit1KeyDown(Sender: TObject; var
Key: Word;
Shift: TShiftState);
begin
if (Key = Key_X) and (Shift = [ssShift]) then
Caption := TimeToStr(Time);
end;
```



Figure 7: Key code constants for use in **OnKeyDown/OnKeyUp** event handlers.

### GAMES

These events work fine for typical X applications, but games require different keyboard handling. When a sprite is moved in response to a key being pressed it needs to move reactively and smoothly. As soon as the key is pressed it should start moving and stay moving until the key is released. The **OnKeyDown/Press** events work via auto-repeat and so only trigger every half second or so, which doesn't do the job.

For cases like these you need a timer triggering its **OnTimer** event very regularly (or else just a loop that keeps executing throughout the game) and there you need to be able to check the state of any given key at that specific moment.

The *Xlib* programming API (accessible through the *Xlib* unit in the *Kylux* RTL) allows you to examine the state of all keys at any given time with a couple of calls. The main one is **XQueryKeymap** but that can only be used after opening a connection to the X server with **XOpenDisplay** (which requires you to call **XCloseDisplay** after).

*Kylux* doesn't come with low level X programming docs but you should be able to find help on these routines with a quick Internet search (one useful *Xlib* documentation site can be found at <http://tronche.com/gui/x/xlib/function-index.html>).

The problem we encountered here (due to limited X knowledge) was the specifics of what **XQueryKeymap** tells you. The docs state it returns an array of bytes (a **TXQueryKeymap**) that identifies which keys are pressed down, each bit representing an individual key. However it says no more on the subject; no mention of how the bits relate to keys on the keyboard.

After some investigation we found that the bit position of each key corresponds to its X key code

(which may differ on different systems). The key code for a given key can be calculated by passing its **keysym** (key symbol) to the **XKeysymToKeycode** X API. A **keysym** is a system-independent way of identifying any key on any keyboard, *Xlib* unit defines constants for all of them. Examples are **XK\_Tab** (the **Tab** key), **XK\_Right** (the right cursor key) and **XK\_F1** (the **F1** key).

The following shows a helper routine from the **XKeyHelper** unit (on this month's disc), which checks the state of a specified key. It takes a **keysym**, which is translated into an appropriate key code for the current X session. The current keymap is then obtained and the key code is used to extract the bit that relates to the requested key. If the bit is set the key is down.

```
uses
Xlib, QForms;

//Pass in a XK_???? constant
function IsKeyDown(KeySym: Cardinal): Boolean;
var
KeyMap: TXQueryKeymap;
KeyCode: Byte;
begin
KeyCode := XKeysymToKeycode(Application.Display,
KeySym);
XQueryKeymap(Application.Display, KeyMap);
Result := Ord(KeyMap[KeyCode div 8]) and (1 shl
(KeyCode mod 8)) > 0;
end;
```

Note that there are no calls to **XOpenDisplay** or **XCloseDisplay** in the code as the *CLX* library makes a connection to the default X server available through its **Display** property. If you need to see the list of available **keysyms**, find them by following these instructions:

- add *Xlib* to your uses clause
- right-click on *Xlib* and choose either Open File at Cursor (or press **Ctrl+Enter**) or Find Declaration and

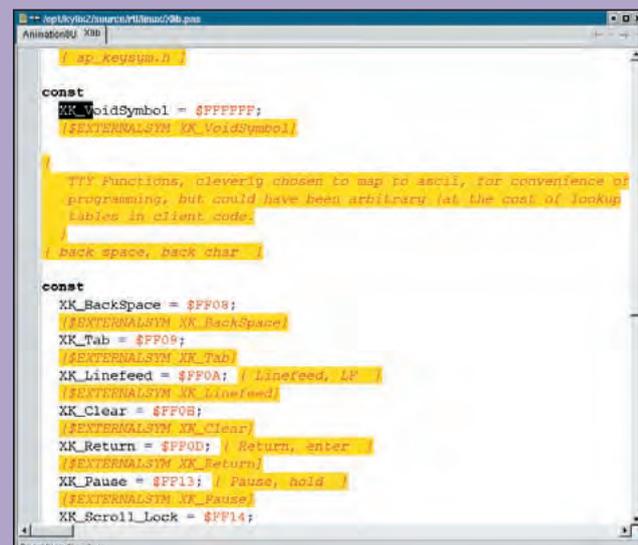


Figure 8: Locating the **keysyms** in the *Xlib* unit.

the *Xlib.pas* unit will be opened in the editor

- choose Search | Incremental Search (or press **Ctrl+E**) and start typing in **XK\_VoidSymbol**

You won't need to type much before the editor locates the first of the **keysyms** as you can see in fig 8.

## TEMPLATE TOOLKITS FOR PERL

# An introduction to CGI programming

**PART 3** Dave Cross rounds off the series looks at some more advanced topics in CGI programming, including Cookies and Perl, and a look at the Template Toolkit.



Whenever you request a page from a website you use a networking protocol called HTTP (Hyper Text Transfer Protocol). This protocol defines the format of the data that your browser sends to the web server and also the data that the server returns to your browser. This protocol was designed to be as simple as possible. One of the downsides of its simplicity is that it is “stateless”.

What this means is that each request to the server is completely independent of any other request. If you request a page from a web site and then click on a link to request another page, the second request has nothing in it to associate it with the first request. Each page request needs to be seen as a separate transaction between the browser and the server.

If you're simply requesting single pages this is not a problem. It can, however, make things difficult when we're trying to write complex CGI programs. For example, imagine a shopping cart program. When visitors come to buy goods from the website, they visit a number of different pages choosing the things they want to buy. When they have finished they expect to see a final shopping basket before they can place their order. Because of the stateless nature of HTTP, the shopping cart program needs to do extra work to keep track of all of this “session” information.

People have invented a number of ways to deal with this problem. Let's take a look at some of the most common.

## Hidden fields

HTML has a type of form input tag which doesn't appear in the displayed page. The tag looks like this:

```
<input type="hidden" name="secret" value="secret data">
```

Although this field isn't displayed, it is passed to the CGI program that processes the form in exactly the same way as visible fields are. This means that you can access the data like this

```
my $secret = param('secret');
```

So if your page is generated by a CGI program and you want to use data from the previous pages but not show it to the user, you can make use of hidden fields. Although hidden fields aren't displayed in the browser, they are part of the HTML for the page so the user can see them by viewing the source of the page.

## Path info

You will sometimes see a URL that looks a bit like `http://www.some.server.com/cgi-bin/program.pl/some/other/data`. On first glance this looks strange. `program.pl` is obviously a CGI prog, but what's that extra data like a file path on the end? This is another part of the CGI specification that we haven't mentioned yet. The extra data is called the “Path Info,” and its contents are passed into the CGI program in an environment variable called

**PATH\_INFO**. In a Perl program that value can be accessed using the variable `$ENV{PATH_INFO}`. `CGI.pm` has a similar function called `path_info()`. Here's a simple program to access the path info:

```
#!/usr/bin/perl -wT

use strict;
use CGI 'standard';

my $info = path_info;

print header,
      start_html,
      h1('Path Info'),
      p('Here are the contents of $ENV{PATH_INFO}');

my @info = split /\/, $info;

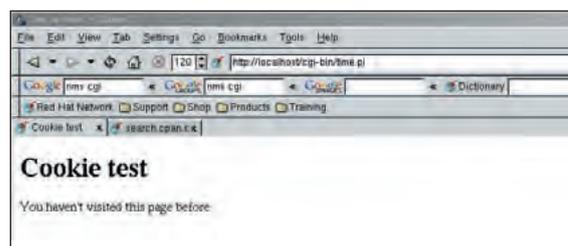
print ul(li(@info));

print end_html;
```

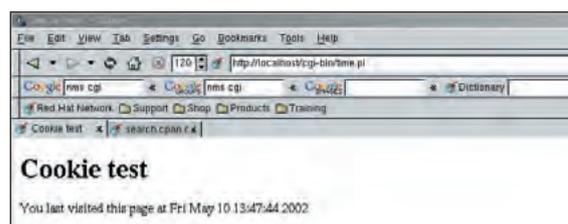
## Cookies

Neither of the previous solutions address the problem of what happens if the user leaves your web site and returns later. For either hidden fields or path info to work, you need an unbroken chain of web pages that you have control of. Netscape

A first visit to the time cookie program.



A subsequent visit to the time cookie program.



came up with a solution to this problem which they called “magic cookies”. Nowadays, they are just called “cookies” but the basic idea has remained the same. Cookies allow a web site to store small pieces of data on the client machine (i.e. at the browser end of the transaction). When the user visits your web site again, the browser will automatically send all cookies belonging to your site back to the server as part of the HTTP header. A cookie therefore has a name and a value (very much like a CGI parameter), a domain that it is associated with and an expiry time. It is also possible to associate a path with a cookie so that it’s only used for a certain area of your site.

CGI.pm contains a function **cookie** which allows you to get and set cookies. Here’s a sample program which uses it.

```
#!/usr/bin/perl -wT
use strict;
use CGI ':standard';
my $time_cookie = cookie(-name=>'time',
                        -value=>scalar localtime,
                        -expires=>'+1y');

print header(-cookie => $time_cookie),
  start_html(-title=>'Cookie test'),
  h1('Cookie test');

if (my $time = cookie('time')) {
  print p("You last visited this page at $time");
} else {
  print p("You haven't visited this page before");
}

print end_html;
```

In this code we deal with two cookies. One that we read from the incoming HTTP header and one that we write to the outgoing HTTP header to be stored by the browser for next time. The code may be confusing because it seems that we process them in the wrong order. We write the outgoing cookie before processing the incoming one. This is because we need to put the outgoing cookie into the HTTP header that we write back to the browser.

Our outgoing cookie has the name **time** and a value that is the current time. We also set it to expire in one year. The **cookie** function returns a string representation of this cookie which we can then use in our call to “header” using the **-cookie** parameter. The browser then takes this cookie and squirrels it away ready to return it to us the next time the user visits our site.

We then look to see whether there was a **time** cookie in the incoming HTTP header. If there is a cookie we extract the value and display it to the user. If our cookie doesn’t exist we display a message saying that the user hasn’t visited this site before.

Most modern browsers allow you to look at your existing set of cookies (sometimes known as a “cookie jar”). In *Galeon*, for example, it is the menu item Tools|Cookie...|View Cookies.

## A more complex cookie example

The most commonly used method for storing session data is using cookies. The big advantage with cookies is that it doesn’t matter whether the user leaves your site and returns later. The data will still be there (unless, of course, the cookie has expired). There are many Internet users who distrust cookies and turn off cookie support in their browsers. See the box *Cookies: The Spy In Your Browser* for more details on why they do this.

Let’s have a look at a more complex example. This program allows a user to log on to your web site and get a personalised

## Cookies

### The Spy In Your Browser

Once Netscape introduced cookies in the mid 1990s, it didn’t take very long for people to realise that you could use them to get more information about users than you could get by tracking HTTP requests. Using cookies allows you to track every visit that a user makes to your site and exactly which pages they look at.

For commercial sites, this is very useful information. For example Amazon use cookies to track which products you are interested in and from that information they can generate recommendations of other products you may like to see.

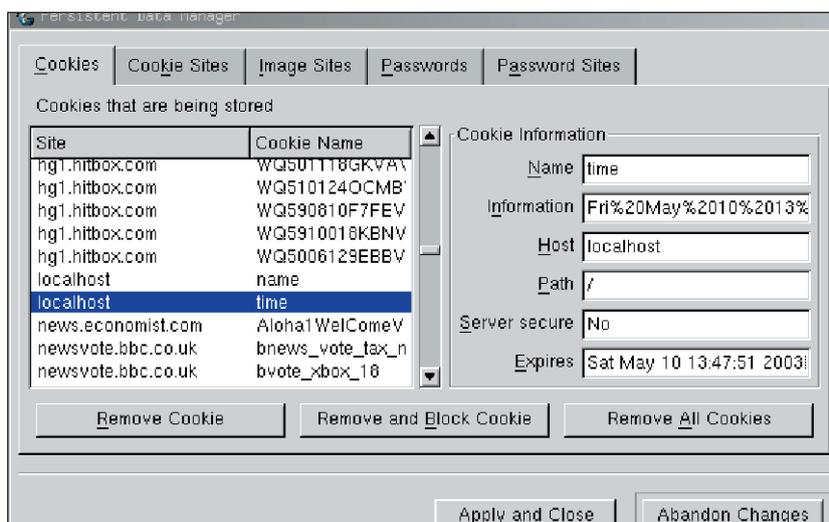
There are a large number of very paranoid people out there on the Internet and when some of them started to see the number of cookies they were building up simply from their day-to-day surfing, they started to get worried. This paranoia is still around today.

If you’ve got a browser that allows you to do so, take a look at the cookies in your cookie jar. You’ll almost certainly find far more cookies than you expected to see. Another interesting

experiment is to spend a day surfing with your browser’s option to warn you before accepting a cookie turned on. You’ll start to get an idea of just how many sites do use cookies to track what you’re doing.

Now I’m not trying to make everyone turn off cookie support in their browser. I’m sure that the vast majority of cookies are benevolent. I like the fact that Amazon tries to guess what other books I’d be interested in. Sometimes they even get it right! I do think, however, that people need to be more aware of what’s going on in their cookie jar. If you can stand it, I recommend always surfing with the “warn before accepting cookie” option turned on. It’s also worth using a browser that has finer grained control over cookie acceptance. For example, many modern browsers allow you to block cookies from certain sites.

A good place to get more information about cookies is <http://www.cookiecentral.com/>. I particularly recommend their FAQ at <http://www.cookiecentral.com/faq/>



version of the site. This is similar to what you’ll see when you visit big sites like Amazon and it welcomes you back by name:

```
#!/usr/bin/perl -wT
use strict;
use CGI ':standard';
my $name;
my $logged;
if (param('login')) {
  $logged = 1;
  $name = param('name');
  print header(-cookie => cookie(-name=>'name',
                                -value=>$name,
                                -expires=>'+1y'));
} elsif (param('logout')) {
  $logged = 0;
  $name = 'Guest';
}
```

**Viewing the data associated with the time cookie.**

The cookie program shows this page when you're not logged in.

```

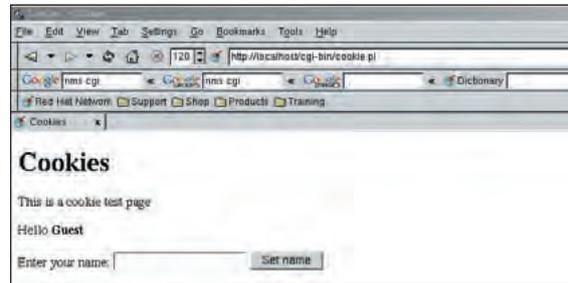
<< print header(-cookie => cookie(-name=>'name',
                                -value=>'',
                                -expires=>'-1d'));
    } else {
        $logged = defined cookie('name');
        $name = cookie('name') || 'Guest';
        print header;
    }
    print start_html(-title => 'Cookies');
    print h1('Cookies');
    print p('This is a cookie test page');
    $name =~ s/</&lt;g;
    $name = b($name);
    print p("Hello $name");
    print start_form;
    if ($logged) {
        print p(submit(-name=>'logout', -value=>'logout'));
    } else {
        print p('Enter your name: ', textfield(-name=>'name'),
                submit(-name=>'login', -value=>'Set name'));
    }
    print end_form;
    print end_html;

```

In this example, we're not only dealing with cookies, but we also have to deal with ordinary CGI parameters, in particular the name of the visitor.

Most of the processing takes place in the three-pronged **if/else** code. If there is a **login** parameter, it means that there is no existing **name** cookie. We set a flag saying that the user is logged in, set the **\$name** variable to the value of the **name** parameter and in the HTTP header we write a cookie containing the name. We set the expiry date (arbitrarily) to be a year in the future.

If there is a **logout** parameter, the user wants our site to forget their name. We set a flag to say that the use isn't logged in



and set the **\$name** variable to the default value of **Guest**. Notice that the **cookie** function doesn't explicitly support the deletion of cookies, so we get around that by setting a cookie that expires one day in the past.

If there are no CGI parameters, then the user is simply visiting our page. We work out whether or not they are logged in from the presence or absence of the **name** cookie and set a flag appropriately. We then print a plain HTTP header as there is no need to alter the cookie.

Having worked out whether or not the user is logged in and what name we should be using, we then display the web page. Depending on the state of the **\$logged** flag, we decide whether to display a login form or a logout button.

## Getting HTML out of your Perl

Most of the web pages that we've created in this series have looked pretty dull. One reason for that is that they were just examples, but another reason is that putting too much complex HTML in your Perl code just makes for over-complex code. Even using the HTML functions within CGI.pm doesn't really help.

Another similar problem is making your CGI-generated pages look the same as the rest of your site. You probably want navigation bars, copyright notices, logos and the like to match your other pages. And if your site design changes, you'll need to change your Perl code to reflect that change. If you're working on a larger site, it's possible that the HTML pages are written by a completely different set of people to the CGI programs.

All of these problems can be solved by removing the HTML from your Perl code, using a templating system. See the box *Templating Systems In Perl* for a review of the systems available and an explanation of why I recommend the Template Toolkit.

We'll reuse the form processing example from two months ago. Here's the new program:

```

#!/usr/bin/perl -Tw
use strict;
use CGI ':standard';
use Template;
my $name = param('name');
my $age = param('age');
my $gender = param('gender');
my @hobbies = param('hobby');
my $list;
if (@hobbies) {
    $list = join ', ', @hobbies;
} else {
    $list = 'None';
}
my $t = Template->new;
print header;
$t->process('form.tt', { name => $name,
                        age => $age,

```

## Template Systems in Perl

### Which module to use?

If you go to <http://search.cpan.org> and do a search for modules with the word "Template" in their name you'll find a huge number of modules. How can anyone choose between them and decide which one to use? The main ones that people seem to talk about are HTML::Template, HTML::EmbPerl, HTML::Mason, Text::Template, Apache::AxKit and the Template Toolkit. Perrin Harkins has written a useful article called "Choosing a Templating System" which you can read at <http://www.perl.com/pub/a/2001/08/21/templating.html>. Here's a brief summary of why I use the Template Toolkit.

Not all of the work I do with Perl is web-based. I like to use a templating system to generate all kinds of data files. Most of the templating systems listed above have been written specifically to be used in generating HTML. You can tell which ones these are as they have "HTML" in their names. In some cases it is possible to use these modules in non-web situations, but it's not really what they're designed for. Apache::Axkit is an even more specialised module. It's for generating HTML

pages from XML and is designed to be used from the Apache web server. It's a great module, but too specialised to be my first choice.

This leaves us with two modules to consider. Text::Template and the Template Toolkit. Both of these are general purpose templating engines. Both allow you to embed processing instructions in a template. The reason why I use the Template Toolkit is that it has its own new language for writing these processing instructions whilst Text::Template uses Perl. Why would I rather not use Perl in my templates? For two reasons. Firstly the Template Toolkit language is specialised for presentation logic. If I was to use Perl for my presentation logic then I might be tempted to do too much processing in the template which might start to erode the distinction that I'm trying to make between processing and presentation. Secondly, it might not be me that is designing the templates and it's far easier for someone to learn the template toolkit language than it is to learn Perl.

So I use the Template Toolkit. You can get more information about it from the website at <http://www.template-toolkit.org/>

```

gender => $gender,
hobbies => $list })
|| die $t->error;

```

As you can see, there is no HTML in the program at all. All of that is handled by the Template module. We load Template.pm at the start of the script and later on we create an instance of the class (Template.pm is an Object Oriented module) using its "new" method. We read the CGI parameters in exactly the same way that we always have and assign them to the same variables.

The actual display of the page is handled by the "process" function which is a method belonging to the Template class. We pass two parameters to the "process" method, the name of a file that contains our template and a reference to a hash which contains the data that we want to display. The "process" method inserts the data into the template and prints the result to STDOUT. If the template processing causes any kind of error then the program dies.

Let's take a look at the contents of form.tt.

```

<html>
<head>
<title>[% name %]</title>
</head>
<body>
<h1>Welcome [% name %]</h1>
<p>Here are your details:</p>
<table>
<tr><td>Name:</td> <td>[% name %]</td></tr>
<tr><td>Age:</td> <td>[% age %]</td></tr>
<tr><td>Gender:</td> <td>[% gender %]</td></tr>
<tr><td>Hobbies:</td> <td>[% hobbies %]</td></tr>
</table>
</body>
</html>

```

As you can see, most of it is just ordinary HTML. The only difference is the presence of [% ... %] tags where you want the data to go. Inside the tags are the names of the data items that you want to display. These correspond to the keys in the hash of data that you passed to the "process" method.

The output that this template produces isn't any prettier than the original form. In fact it's exactly the same. But I hope you can see that, by separating the processing from the presentation like this, it's made things easier to handle. Let's make some changes:

```

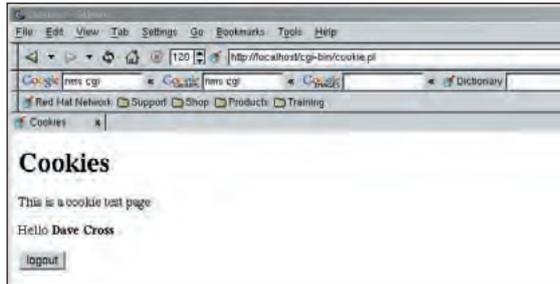
#!/usr/bin/perl -Tw

use strict;
use CGI 'standard';
use Template;

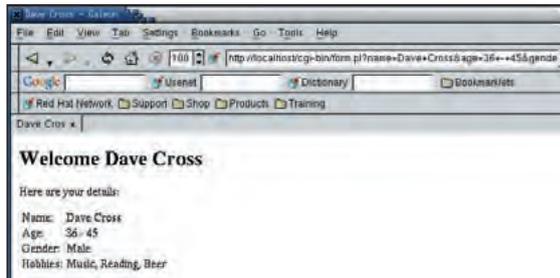
my $name = param('name');
my $age = param('age');
my $gender = param('gender');
my @hobbies = param('hobby');
my $t = Template->new;
print header;
$t->process('form2.tt', { name => $name,
                        age => $age,
                        gender => $gender,
                        hobbies => \@hobbies })
|| die $t->error;

```

The only change we've made is that we no longer check on the number of hobbies and create a text list in the program. We now pass a reference to the list to the template processor.



The cookie program shows this page when you are logged in.



Standard CGI Program Output.

We'll also change the template.

```

<html>
<head>
<title>[% name %]</title>
</head>
<body>
<h1>Welcome [% name %]</h1>
<p>Here are your details:</p>
<table>
<tr><td>Name:</td> <td>[% name %]</td></tr>
<tr><td>Age:</td> <td>[% age %]</td></tr>
<tr><td>Gender:</td> <td>[% gender %]</td></tr>
<tr><td>Hobbies:</td> <td>
[% IF hobbies.size %]
<ul>
[% FOREACH hobby = hobbies %]
<li>[% hobby %]</li>
[% END %]
</ul>
[% ELSE %]
None
[% END %]
</td></tr>
</table>
</body>
</html>

```

We've added some logic to the template using simple conditional and looping constructs. As the **hobbies** template variable is now a list of items we can check the number of items in the list using the **.size** method. If the list is empty we display the word **None**, but if the list contains elements we create an unordered list of the items.

## Conclusion

So that concludes our three-month introduction to CGI programming. I hope it's started you thinking about the multitude of things that you can achieve using this technology. If you'd like to learn more, the two best books on the subject are probably *CGI Programming with Perl* by Guerlich, Gundavaram and Birznies and *Writing CGI Applications with Perl* by Meltzer and Michalski. [LXF](#)

# 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 modern woes to network administrations, we can find the answer for you – just fire off a letter or email and it'll all be taken care of.

LXF answers guy

### David Coulson

is a networking and security guru with plenty of sysadmin experience to boot.



### 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*...



## Mute Mutt keys

**Q** I want to read and manage my messages in *Mutt* with the following command:

```
mutt -f pop://<account>@<pop server>.<ISP>.<whatever>
```

How do I do this, *i.e.* get into pager while on a remote pop server? The commands listed are **q:Quit**, **d:Del**, **u:Undel**, **s:Save**, **m:Mail** (*i.e.* new message), **r:Reply** (as in reply to message), **g:Group**, **?:Help**. I know this must be very basic, but I don't get it. I have looked at the manual and all sorts of other resources, but either this issue isn't listed or I am suffering from info overload. Can you help? I use SuSE 7.3 and *Mutt*1.3.22.1i.

**Brian Durant**

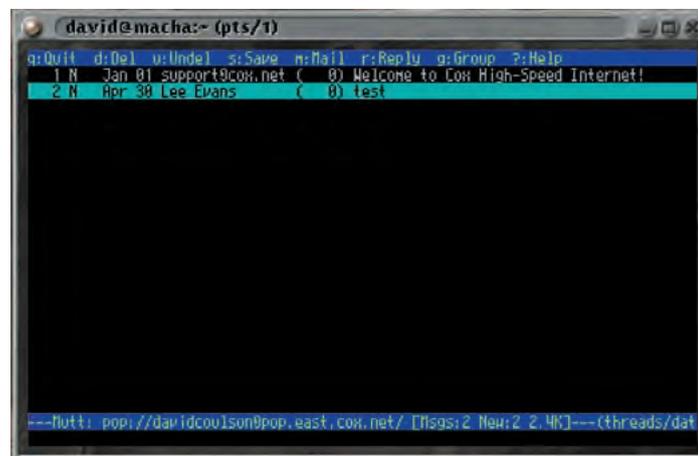
**A** We're not 100% sure that we know what you're asking. In *Mutt*, the pager is the section of the program which displays emails. To have *Mutt* display the selected email in the index list, you simply need to press **Enter** when the appropriate email is selected. If you're trying to find something else, pressing **?** will list all of the key bindings which are available to *Mutt*, which, while being rather cryptic, should point you in the appropriate direction.

## Mandrake install

**Q** I'd like to install Mandrake 8.2 from this month's disc onto my new PC, which has XP home edition on it. Have you guys tried this? Do you know if there are any known problems with it? Any advice about recording hardware details before attempting an installation?

**Neil Earnshaw**

**A** Mandrake 8.2 has many bug fixes, but little of the actual structure of the distribution has changed, unlike when we went from 7.x to 8.x. The easiest way to



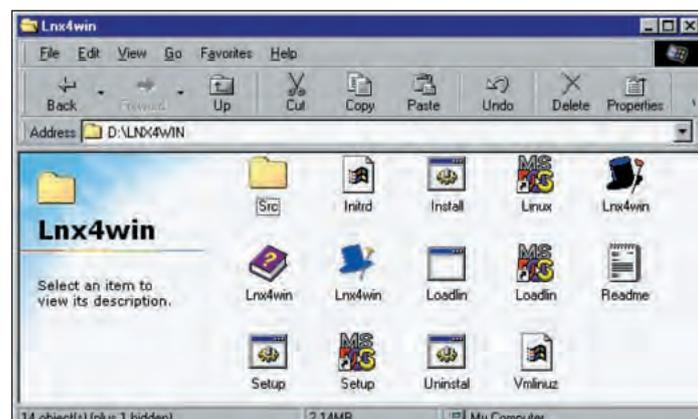
The mail client *Mutt* supports connections to most common server types, including POP3 and IMAP4.

record hardware information is to use the device manager in Windows to record the IRQ and other details of specific pieces of hardware, such as the video card, sound card, NIC if any, and so forth. Most hardware now is Plug-and-Play, so it's unlikely that you will need to do anything with IRQs and I/O addresses yourself. Knowing the model and manufacturer of the chipsets on the main hardware is probably all you need at first, and even then, *XFree86* is knowledgeable enough to auto-detect the video card. Many people use Mandrake, so if

you're worried about problems with it, then it may be worth looking at the [www.linux-mandrake.com/en](http://www.linux-mandrake.com/en) website and seeing what updates are available for their 8.2 release. You may also want to post on the LXF Forums, to find out if other users who have installed it from the LXF coverdisc have experienced problems with any specific hardware.

## OS within OS

**Q** According to the install help file on your latest disc, the Mandrake 8.2 one, you can



Lin4Win from Mandrake allows you to install your Linux system within Windows, rather than repartitioning.

install Linux in Windows by booting from CD, pressing **F1** and typing **lnx4win**. I tried this but it says **lnx4win** is not there, any ideas?  
**John Marsh**

**A** Installing Mandrake using *lnx4win* installs Linux in an image file on the Windows partition, rather than re-partitioning the drive and creating swap and filesystem partitions for Linux. Unlike many other distributions, Mandrake installs Linux on an *ext2* filesystem, rather than UMSDOS, so it's a little bit quicker.

If you're going to use *lnx4win*, then you will want to boot up into Windows, then run *lnx4win.exe* within the LNX4WIN directory on the CD. Booting from the CD is only necessary if you're going to do a 'proper' installation of Linux, rather than one from within Windows.

## IrDA and dialup

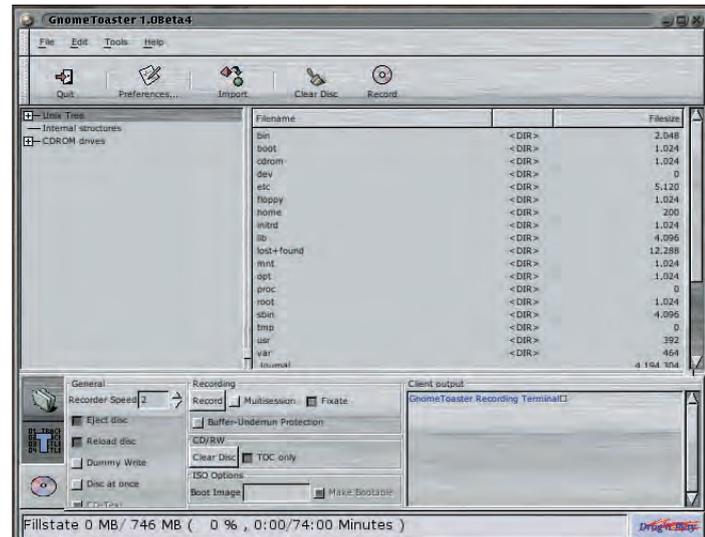
**Q1:** In the PC I have a CD-RW. The line in *fstab* that mounts it is shown below.

```
/dev/scd0 /mnt/cdrw auto user,
iocharset=iso8859-1,umask=0,exec,
codepage=850,ro,noauto 0 0
```

The line for the CD-RW has **ro** near the end of the line. I understand that means read-only. So how am I able to write a CD-R or CD-RW in the drive?

**Q2:** The powerbook has an infrared device. I installed an infrared device on the PC and connected it to the mainboard which has connector and "hardware". I have read the HOWTOs but am left with the idea that instead of configuring I must write a program. I am not a programmer. How and what do I configure to get Mdk and YDL talking to each other like *ftp* or *telnet*, over the IrDA? Or where can I find a clearly explained HOWTO?

**Q3:** My PC and G3 are connected as a LAN using ethernet and cross over cable. I can *ping*, *ftp* and *telnet* over my LAN. If I dial out to my ISP with the powerbook it appears not to connect until I checked the dial configurator which said I was



Burning a CD is done with a program such as *cdrecord*, which is the underlying program used by the GUI, *Gtoaster*.

connected. I tried a *ping* and the report said my IP was 192.168.1.2. Shown below:

```
ping 216.22.22.22 (216.22.22.22)
from 192.168.1.2 : 56(84) bytes of
data.
```

After doing *ifconfig* one or two times, I would get a report that my IP was 66.66.66.66 as below:

```
ppp0-inet addr:66.66.66.66
P-t-P:66.66.66.67
```

For the rest of that session every thing would work OK until a new connection. The dialup functioned properly several times before doing this. How do I make the dialup work correctly? **Note:** The PC did the same thing. In trying to fix the problem on the PC, it would hang on boot and I had to reinstall Mdk 8.1. The kind of thing one has to do with Windows OSs. **Q4:** I would like to "back up" a directory (not the whole disk) that contains directories and files to a CD-RW. I want to copy the directory to the CD-RW and have only the newly added files get copied (backed up), not the whole directory each time. I found a command line program I think will do this (*cdbackup*) but haven't found a GUI program. Any suggestions?  
**Mike**

**A** Answering in order.  
1/ You don't write to a CD-RW by mounting it, so the way it is mounted has no bearing on the ability to write to it. You need a program such as *cdrecord*, which writes directly to the CD-RW. It's important that the CD is not mounted when you try to write to it, otherwise you'll experience a few problems. If you're after a GUI solution, then there are many front-ends for *cdrecord* which let you drag and drop files for burning. Check

## A QUICK REFERENCE TO: email

One you're connected to the Internet, it's rather useful to be able to collect your email and read it. Many email clients for Linux can handle POP3 and IMAP4 mail servers themselves, so you can put in the hostname of the incoming mail server, and then have it download itself. However, many clients, including *Pine* and *Mutt*, allow you to read email from a file located on the local machine, avoiding the need to be connected to the Internet when you do your email. This needs the system to have a mail server such as *Sendmail*, *exim* or *postfix* installed, so it can handle the email sent from the mail client while the system is offline.

If you don't have a mail server installed at the moment, then either *postfix* or *exim* are the best choices for a beginner, as they are far easier to install than *Sendmail*. In either case, you will need to ensure that the mail server is not an open relay, and will only accept email messages from 'localhost' or '127.0.0.1',

otherwise people can use your machine to distribute UBE/UCE and other types of spam over the Internet. You can test your mail server by using the *mail* command to send yourself an email;

```
$ mail david
Subject: Test
Example
.
```

The mail will usually end up in */var/spool/mail/\$USER*, or whatever the *\$MAIL* variable is set to. We want our email to be delivered to this mailbox, so we can access it using *Mutt* or *Pine*, and we need to use a separate program, *fetchmail*, to collect our email and deliver it to the correct place on our system.

*fetchmail* consists of a single configuration file, *.fetchmailrc*, which contains information as to which server to use, the username and password, how to deliver it locally, and so forth. We might use an entry in *.fetchmailrc* similar to; `poll pop3.demon.co.uk proto pop3 user dcoutson pass PASSWORD to david`



This will connect to `pop3.demon.co.uk`, using the POP3 protocol, and collect mail for the user 'dcoutson' and deliver it to 'david' locally. *fetchmail* can have as many entries as you like, so if you have mail on two or three different providers, then you can have it all delivered to the same place.

For outgoing email, you will need to enable *smarthost* relaying in your mail server, so that email is sent from your mail server to your ISP's SMTP server, rather than straight to the destination, as it is difficult to deliver email everywhere if you have a dynamic IP address due to spam filters.

# FREQUENTLY ASKED QUESTIONS: FIREWALLS

## FAQ How to I stop remote machines connecting to my computer?

There are two different ways to block access by remote machines to your computer. The first is to restrict access within particular services, so that they deny access to specific hosts. Many use `/etc/hosts.allow` and `/etc/hosts.deny`, and you might set yours up as:

```
hosts.allow;
```

```
ALL: LOCAL
ALL: 192.168.
```

```
hosts.allow;
```

```
ALL: ALL
```

This will block access to all machines, except those with an IP starting 192.168, or if the connection is local to the machine. However, many services don't take advantage of `hosts.allow` and `hosts.deny`, so we need to use something else.

Depending upon your Linux kernel version, you can either use `ipchains` or `iptables`. Linux 2.4 users

should use `iptables`, as it has far more advanced firewalling capabilities than `ipchains`. We can block access to our machine on a specific interface, such as `ppp0`, using the following:

```
iptables -A INPUT -i ppp0 -j DROP
```

However, this will drop all connections, including those which did not originate externally, so we need to allow connections which we started to work:

```
iptables -I INPUT -i ppp0 -m state --state RELATED,ESTABLISHED -j ALLOW
```

We can check our firewall rules with:

```
iptables -nvL
```

`iptables` requires kernel modules to be compiled, and a user-space application, `iptables`, to be compiled. The latest release of `iptables` can be found at [www.netfilter.org](http://www.netfilter.org).

## FAQ Can I permit access to specific ports?

Using `iptables`, we can allow packets belonging to a specific IP protocol through, such as UDP or TCP. It's important that we allow packets

before they are dropped, so we have to use `-I` to insert it before our **DROP** rule:

```
iptables -I INPUT -p tcp --dport 22 -j ACCEPT
```

which would allow access to the SSH port, 22, from anywhere.

## FAQ How can I redirect a port to a box on my network?

In order to redirect a port, we have to use the DNAT, or destination NAT, capability of `iptables`. This rewrites a packet so that the destination IP of the packet changes from the IP of our external interface IP, to an IP of an internal machine, so the packet is routed to it. If we wanted port 22 on the gateway connected to the Internet to be redirected to a machine on 10.11.2, we could do:

```
iptables -t nat -A PREROUTING -i ppp0 -p tcp --dport 22 -j DNAT --to 10.11.2:22
```

It is perfectly possible to redirect one port onto another, so rather than being redirect to port 22, we could redirect it to 10.11.2:3452 if we so desired.

## FAQ Can I share my Internet connection with other systems on my network?

In order to share the network connection, we need to use IP Masquerading. Assuming the network is set up correctly, and the IP of our machine is the gateway for the other machines connected to the network, we only have to enable IP forwarding, then set up masquerading. To enable IP forwarding, we just need to do:

```
echo 1 > /proc/sys/net/ipv4/ip_forward
```

And for masquerading, assuming `ppp0` is the interface connected to the Internet:

```
iptables -t nat -A POSTROUTING -o ppp0 -j MASQUERADE
```

## FAQ What's the difference between SNAT and MASQUERADE?

If we have a connection on a dynamic IP, then we need to use the **MASQUERADE** rule, as when the connection goes down, the connection tracking system will

freshmeat.net for those.

2/ There are IrDA modules in the kernel, which support various forms of IrDA data, including IrLAN, which is an IP network over IrDA. To enable IrLAN, one would do something similar to:

```
irattach /dev/ttyS0
```

```
irmanager -d 1
```

```
insmod irlan access=2
```

```
ifconfig irlan0 10.0.0.1
```

```
netmask 255.255.255.252 up
```

On the other machine, do the same thing, but use **10.0.0.2** instead. Note that you will need to change the **irattach** line appropriately to point to

the correct serial device to which the IrDA hardware is connected. There are IrDA management tools, which should be available with your distribution, so that you can do this all automatically when you boot up.

3/ As you've not provided a great deal of configuration information, it's somewhat difficult to come up with a perfect solution to this problem. You did not say what you were doing with `ifconfig`, but we would assume that it is a routing problem, and you've got the default route for packets setup to be the LAN, rather than the dialup connection. You need to pick one machine to use for outgoing dialups, then remove the default route from the network configuration, and on the other machine, set the gateway to the 192.168.0.x IP of the machine doing the dialup. If you're still having problems, it may be worth looking at the routing tables with `route -n` and finding out where packets are being sent. Also remember to setup IP

Masquerading on the machine doing the dialup, although it is best to check that the box with the modem can actually access the Internet first. 4/ `cdbackup` does exactly what you need, so what is wrong with it being a console application?

## Moving snapshots

**Q** In LXF27 the question was asked how to capture screenshots including programs like the *Ogle DVD player*.

The answer is to use KDE's *Ksnapshot* which will happily grab screen images including DVD output. It can capture either the whole desktop or a single selected window and exports the image in png format. It works perfectly in GNOME also.

Enjoy!

Andrew Richardson

**A** From our own testing, it would appear that *Ksnapshot* cannot capture Xvideo output as the

```
david@macha:~ (pts/1)
port port_number
The port option sets the I/O port, as described
above.

irq irq_number
The irq option sets the hardware IRQ, as described
above.

uart uart_type
This option is used to set the UART type. The per-
mitted types are none, 8250, 16450, 16550, 16550A,
16650, 16650U2, 16654, 16750, 16850, 16950, and
16954. Using UART type none will disable the port.

Some internal modems are billed as having a "16550A
UART with a 1k buffer". This is a lie. They do
not really have a 16550A compatible UART;
instead what they have is a 16450 compatible UART
with a 1k receive buffer to prevent receiver over-
runs. This is important, because they do not have
a transmit FIFO. Hence, they are not compatible
with a 16550A UART and the autoconfiguration pro-
gram.

Manual page setserial(8) line 441
```

The `setserial` utility can be used to change the IRQ and IO port of serial devices such as modems.

automatically flush all of the rules belonging to the interface as a security measure. However, if we have a static IP, this will stop TCP connections from resuming, which is rather annoying. By using a SNAT rule, we can have all the features of the **MASQUERADE** rule, but if the connection goes down, tracked connections will remain in the list.

### FAQ Can I log unauthorised connections to my computer?

You can log packets through *syslogd* using the **LOG** rule in *iptables*. Usually, two separate rules are used for each set of packets to be dropped, one for the logging, and the other to drop the packets. To log all incoming connections, then drop them, we would do;

```
iptables -A INPUT -i ppp0 -m state --state NEW -m limit --limit 1/sec -j LOG
iptables -A INPUT -i ppp0 -m state --state NEW -j DROP
```

By specifying a log limit, we can avoid attacks which fill /var.

earlier reader wanted. We would guess that your video card does not support the *Xv* extension, so it is using *Xshm* or *X11* rendering, which can be captured by a program such as *Ksnapshot*. Many applications, including *Xine* and *Sinek*, have a button to dump the current video frame to disk, so when you make a screenshot, it is possible to paste the captured frame over the blue section of the screenshot. It is by no means perfect, but it is a better solution than using *X11* to watch a DVD, which is awfully slow, even on a very fast machine.

## Partitioning

**Q** *LXF1* inspired me to try Linux and I've been using it and reading *LXF* ever since.

However, I have now upgraded my system by purchasing an Evesham Axis 1900+RK system. It has 512MB RAM and an 80GB hard drive running Windows XP

on a single 80GB Partition.

I need a dual boot system but run most day-to-day work on Linux – but I am unable to non-destructively re-partition the Windows XP partition in order to load Red Hat 7.2 Linux, which is the distribution I use.

*DiskDruid* says it is unable to make a partition and *fips* does not recognise the filesystem.

Can you tell me how to non-destructively repartition a Windows XP partition so that I can load Red Hat 7.2 Linux?

**Gus Hills**

**A** Unfortunately, you will need a program such as *Partition Magic* in order to resize the NTFS partition used by Windows XP. As far as we know, no distributions of Linux can handle resizing NTFS filesystems, so you're going to have to spend some money to have it done properly. An alternative may be to use a distribution such as Mandrake which has an easy method to install Linux into the Windows partition, so you need not mess around with partitions.

## Serial IRQs

**Q** I have loaded the Mandrake 8.2 distro from the May issue of *Linux Format*. It is great. However, I am having difficulty setting up my ISA modem.

The modem is on COM port 3 and requires the use of IRQ 5. Please please tell me how to set COM port 3 to use IRQ 5. So far I have spent about 12 hours rooting through manuals, help pages and HOWTOs and I have got nowhere. I am relatively new to Linux, but generally considered no fool. It does seem to me that the biggest failing (if that's the right word) of Linux is that there are no step-by-step guides for the many necessary operations in setting up a working system. I do feel that an 'Idiot's Guide' to obtaining a very basic installation would give many more potential Linux users the confidence to press on to a full working system.

My sincere thanks for a great magazine I enjoy the read even though a lot of it is over my head. Many thanks.

**Sandy**

**A** *setserial* can be used to set the IRQ for modems. Windows COM 3 is ttyS2 in Linux, so we can set the IRQ to 5:

```
setserial /dev/ttyS2 irq 5
```

It's also important to set the port, which is the bottom I/O address in the card's range. As the card is ISA, it may be worth using *isapnp* to extract all of this information from the card's configuration.

## WPO log writing

**Q** I am trying to get *WPO2000* to run on SuSE 7.3.

However, I cannot get the program to run, because of errors as stated below.

```
johnbywr@(none):-> wordperfect
```

```
Notice: Using config file:
```

```
/etc/fontastic/config
```

```
Error: Cannot open error file
```

```
"/var/log/fontastic/fs-errors"
```

```
Notice: Font server started:
```

```
Wed May 1 12:09:12 2002
```

```
Copyright 2000 Bitstream Inc.
```

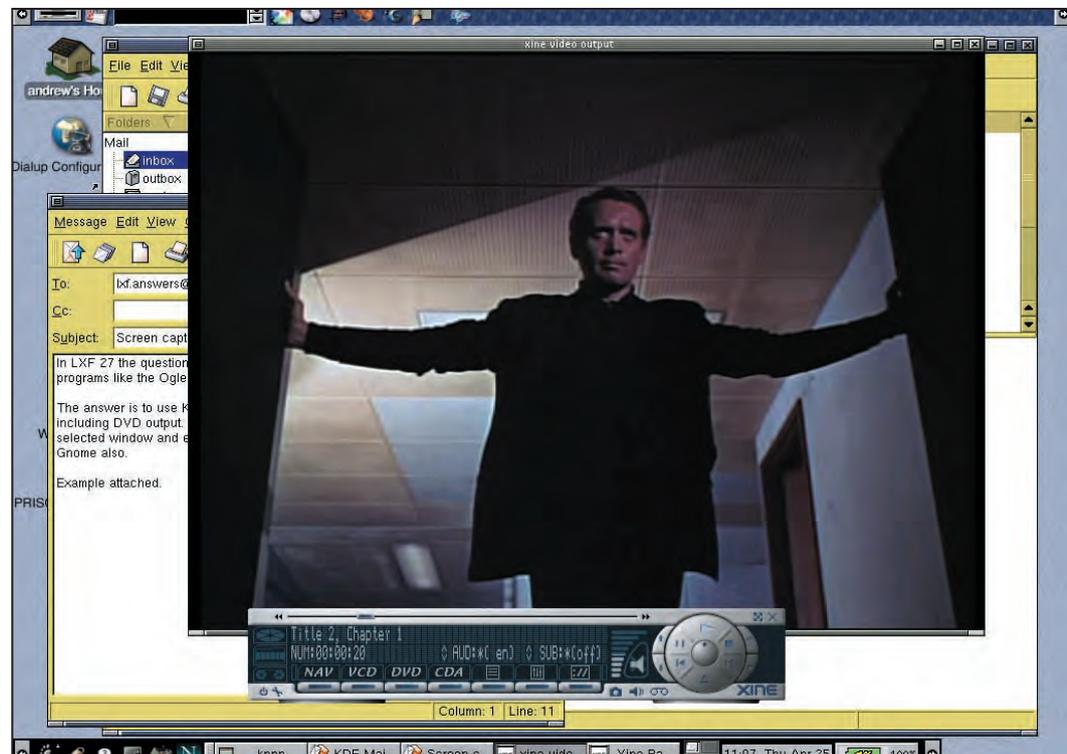
```
SDK Version 1.3.6
```

Then a box comes up with: "No running font server detected. I will try to start a font server on this host and wait for 30 sec to allow the font server to start up."

Then another box stating:- "unable to add Fonttastic font to the font path. The font server is probably not installed or not running. Correct the problem and try again."

I have been on both SuSE's and Corel's sites, without success in finding a solution. Can the Fonttastic fonts be added to the path, or a symbolic link be made?

I installed *WPO2000* with the script from Corel. Which I had to change slightly as SuSE uses *libc 2.2*, and *WPO2000* wanted only *libc 2.0* or *2.1*.



*Ksnapshot* can capture screenshots of DVDs played using *X11*, but if they use *Xv* you will just get a blue box.

« I have never had trouble like this using *WPO* for windows. Which is one of the reasons people are put off trying Linux. I find that generally *SuSE* is very good, but then problems like this arise. And your ability is suddenly shown to be so lacking.

I do hope that you can assist in sorting this problem, as with *WPO* I will feel much more at home.

John Bywater

**A** *Fonttastic* can't write to the error file, as `/var/log` is only writable by root, and you're running it as 'johnbywr'. If you create the `/var/log/fonttastic` directory, then:

```
chmod 777 /var/log/fonttastic
```

so that everyone can write to it, you will be able to check up on the *fonttastic* errors to find out what the problem is. Hopefully, this will explain why *WordPerfect* does not think it is running, as it may not be able to bind to a port, or read it's configuration file when being run as a non-root user.

Depending how your X system is setup, you should be able to add the font directories to `/etc/X11/xfs/config`, or to your `/etc/X11/XF86Config`.

## Vaio not viable?

**Q** I have a Sony Vaio SR27, yes one of those mini notebooks with no Parallel or serial port, no floppy disk and a CD connecting via a PCMCIA card.

I had Mandrake 7.2 running, dual booted with the pre-installed Win ME. Then I received this month's CDs and decided it was time to upgrade to 8.2.

Unfortunately it didn't work.

I got it to boot from the CD using **F1** then typing

```
expert all ide1=0x180, 0x386
```

and all looked good until I got to the section following 'format partitions' which should be package selection. At that point, having reformatted my Linux partition, I got the message "An error occurred missing base system package" and a push button "OK?" On pressing the push button it returned me to the partition selection screen.

Linux has been running well on my laptop and I have encouraged many others to switch from Windows or at least try it out. This has been my first let down. Whilst I accept that you do not produce Mandrake, any clue to this problem would be greatly appreciated.

Regards from an avid reader,  
Colin Thomas

**A** You should be able to switch to a console using **Ctrl-Alt-F2**, where you can check that Mandrake found your CD drive and IDE controller and that it was able to mount it properly. An option may be to copy the contents of the Mandrake CD to the Windows section of the hard disk, then do an install from there, rather than from the CD.

As you managed to get Mandrake 7.2 installed, it may be worth checking to ensure that you can still install that release of the distribution before trying again with 8.2. The SR27 is a common laptop, and many people have installed Linux on it, looking at [www.linux-laptop.net](http://www.linux-laptop.net), so it may be worth looking at the site and checking to see how other people have installed Linux on their Vaios, even if the model is not exactly the same. You might get lucky, and someone else may have had problems installing Mandrake 8.x using the method you are using.

## KDE RPMs

**Q** I have Red Hat Linux 7.0 running with *KDE1* quite happily (the PC multi-boots Linux, Win2000 and Win98!)

I'm trying to install *KDE 2.21* from the Dec 2001 coverdisc (CD) and I'm hoping you can help me with some issues. Ultimately I hope to install the office system from the Nov 2001 coverdisc. Here's what I've done...

1. *tar*, on my system, did not have the switches (`-f`, I think) so I managed to get a later version of *tar* and install it. So far so good!
2. Tried building the first KDE package (*kdelib* I think, sorry about the vagueness, I'm at work right now) but it said that I needed a later *Qt*. So I downloaded that from Red Hat
3. Tried to install *qt-2.3.0-3.i386* as follows...

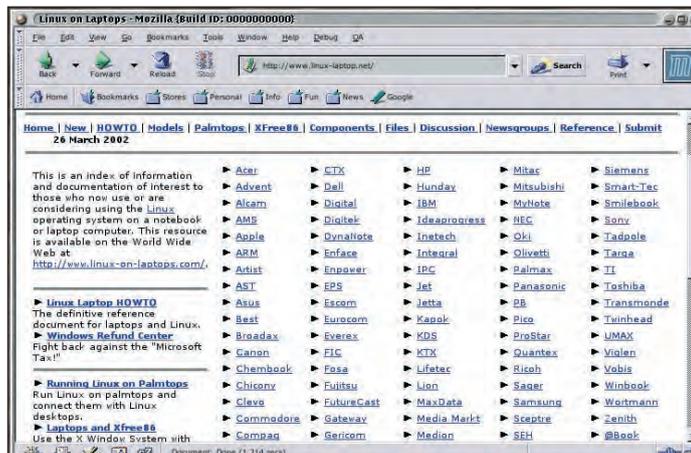
```
rpm -Uvh qt-2.3.0-3.i386.rpm
but it says it requires the following..
libXft.so.1
```

```
libXrender.so.1
```

```
libfreetype.so.6
```

```
libmng.so.1
```

I'm not sure where or how to get these. When I search the web for them it seems to mostly talk about *XFree86-libs-4.???.i386* and Red Hat 7.1 features



The Linux Laptop site ([www.linux-laptop.net](http://www.linux-laptop.net)) has configuration and installation information for nearly all laptop brands you can think of.

As you might understand this is all a bit bewildering but probably because I'm new and need to read more.

Hope you can help with these questions about the approach from here...

**Q1.** Should I persist with getting the newer *XFree86 rpm* and if so can you tell me a place to get it from ?

**Q2.** Will this just uncover more dependencies? If so, is there a way of seeing ahead of time what components will need to be upgraded so I can get them all at once? *E.g.* I don't even know what *libXft* is or what version is running on my system.

**Q3.** Would I be better off just upgrading my whole system to the latest Red Hat (7.3?) to get all of the latest things on one go?

Rod Burke

**A** If you're wanting to install *KDE 2*, then it may be best to install it using RPMs specifically built for Red Hat 7.0, as if you use RPMs for Red Hat 7.2, then it will expect you to be running *XFree86 4.x*, rather than 3.3.6. You should be able to install *XFree86 4.2.0* using RPMs without too much trouble, although the RPMs are quite large to download. Unless you download RPMs for a different version of the distribution than you are running then you should not experience any dependency problems. Red Hat 7.0 is generally considered to be quite problematic, as it was the first release of the 7.x series of the distribution from Red Hat. Upgrading to 7.3 will fix all of the problems you are experiencing, but you may also notice that quite a few things break and get

confused, particularly if you've been building packages from source in the past, as the dependencies will not match up with what has been installed using RPMs. All of the libraries which *Qt* is complaining about, apart from *FreeType*, are distributed as part of *XFree86 4.x*, so simply upgrading *XFree86* will fix that specific problem. *XFree86* should be able to be found at [www.rpmfind.net](http://www.rpmfind.net), or on Red Hat's FTP server at <ftp://ftp.redhat.com/> **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: [linuxformat@futurenet.co.uk](mailto:linuxformat@futurenet.co.uk)

# Coverdisc

Neil Bothwick is your guide through the wonders of this month's jam-packed Linux Format CD. Can you contain your excitement?



## On the CD

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

### Essential info

On page 105 we have grouped together essential info on the different types of packages on your coverdiscs – along with instructions for installing source packages.

### Important notice

Before you even put the CD or DVD in your drive, please make sure you read, understand and agree to the following: The Linux Format CD/DVD 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/DVD 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.

## READ ME FIRST

After being in beta for a long time, *OpenOffice.org* has now reached version 1.0 and is here on the CD and DVD. It's a browser bonanza this month, as we bring you *Opera 6* for Linux, *Mozilla 1rc3* and *Netscape 7pr*. You can install *Netscape 7 pre-release* direct from the CD or DVD.

When the installer starts, the browser will prompt you to download the latest version from

*Open Office 1.0* is out, and on your Linux Format coverdiscs.



Netscape. There may well be a new version by the time you read this, but it may be worth

checking the numbers first, because at the moment this message always appears!

## OpenOffice.org

It seems as though *OpenOffice.org* has been around in beta form forever, but it has finally reached full release status. We have version 1.0 of Sun's open source, integrated office suite on the CD and DVD for you. The Office/OpenOffice directory contains four archives and some text files. For a standard installation you will need one of these archives. 'Either `OOo_1.0.0_LinuxIntel_install.tar.gz` if you have an Intel/AMD CPU or `'OOo_1.0.0_LinuxPPC_install_en-US.tar.gz` if you use a PowerPC.

Users of other hardware architectures will have to compile from the source code, which is in `'OOo_1.0.0_source.tar.gz`'. There is no README file in the root of this archive, but compilation instructions are in the other two archives (no, I don't understand that either).

For now, we will assume you are installing from one of the two binary archives. Installation takes place as a normal user, you don't need to be root for this. While in your home directory, type

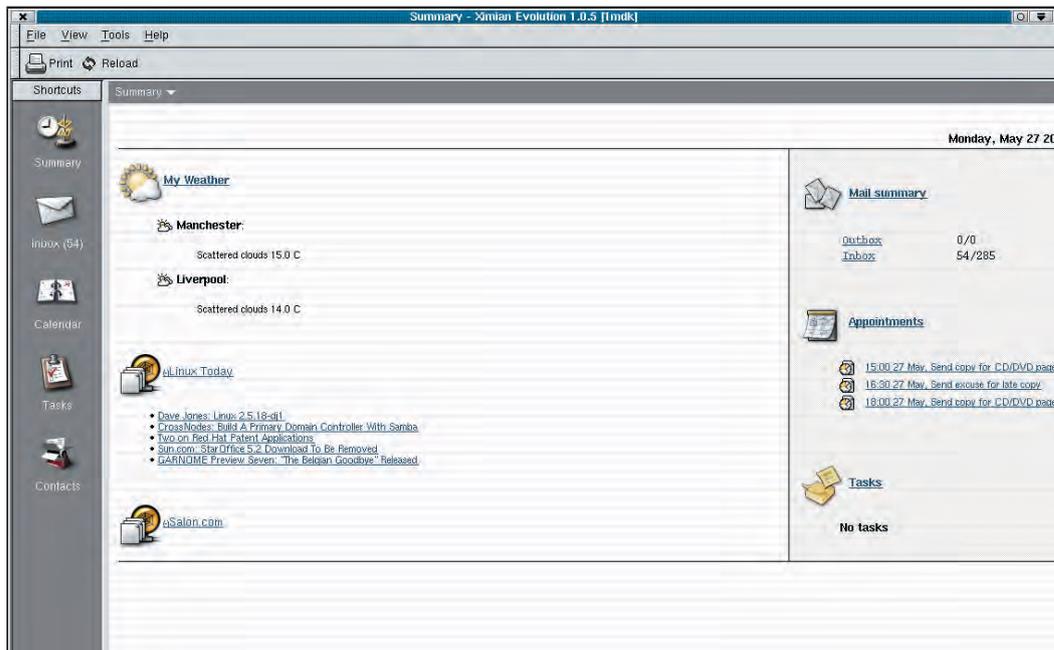


Installing Open Office is a breeze, unpack the archive and type `./setup`.

```
tar xvf /mnt/cdrom/
Office/OpenOffice/
OOo_1.0.0_LinuxIntel_install.tar.gz
changing the path to your CDRROM
drive and the name of the archive if
necessary. Now type
cd install
./setup
```

to start the installer. Installation is very straightforward. You will get a message if it cannot find Java, but it works

without Java. The installer creates a directory called `OpenOffice.org1.0`, assuming you accepted the default, and puts everything in here. You may now delete the install directory. To start *OpenOffice.org* type `~/OpenOffice.org 1.0/soffice` or you can add this command to your KDE or GNOME menus. The first time you run *OpenOffice.org* (*OOo*) it asks for some information. The usual name,



### Manage your email, your contacts, most of your life, with Evolution.

address, company, telephone etc. It doesn't appear to send this information anyway when you give it (I wasn't running any sort of network monitor at the time but the LEDs on the hub didn't flash) but you will be offered the chance to register with Sun the next time you run the program. If you feel uncomfortable with giving out this information, either leave the fields blank or fill them with garbage, *OpenOffice.org* should still work.

When you start *OOo*, it loads up as word processor, you then use the menus to change to a different format if that's what you need. For a quicker way, install *OpenOffice.orgQuickStarter (oooqs)*. This puts an icon in the KDE panel that lets you start up *OOo* in whichever mode you need. There

is extensive online help in *OOo*, so install it, start it up and take it for a spin round the block.

## Evolution

*Evolution* is described as "Integrated Workgroup and Personal Information Management for Linux and UNIX". In other words, it combines email, contacts and appointments in one package.

Installing *Evolution* requires that some dependencies are satisfied first. These are all in the Internet/Evolution directory of the CD and DVD and need to be installed in the correct order. We have packages for the main distributions as well as source code tarballs, so use whichever is most suitable.

*bonobo-conf*  
*gal*  
*gtkhtml*  
*evolution*

If working from tarballs, you should unpack each archive and do the standard **./configure && make && make install** on each. Mandrake, Red Hat or SuSE users should **cd** to the relevant directory of the disc and type **rpm -Uhv \*.rpm**

Once installed, you can use *Evolution* to handle your email, diary, TODO lists and contacts. It can even tell you what the weather's like outside. This is handy if you work for a slave driver who won't even let you take time to look out of the window ("That's enough whinging, Bothwick. Get back to compiling that *coverdisc!*" – ed).

One feature on its own is enough to endear *Evolution* to me, the ability to hide previously read emails. I keep a lot of old emails and don't like the way many other email programs leave them cluttering up the display, although calling *Evolution* an email programs does it an injustice, it's far more than that.

To install from source, you need the GNOME 1.4 development libraries, these should be on your distribution's CDs.

The packages you need to install are, in order:

## Multiple versions?

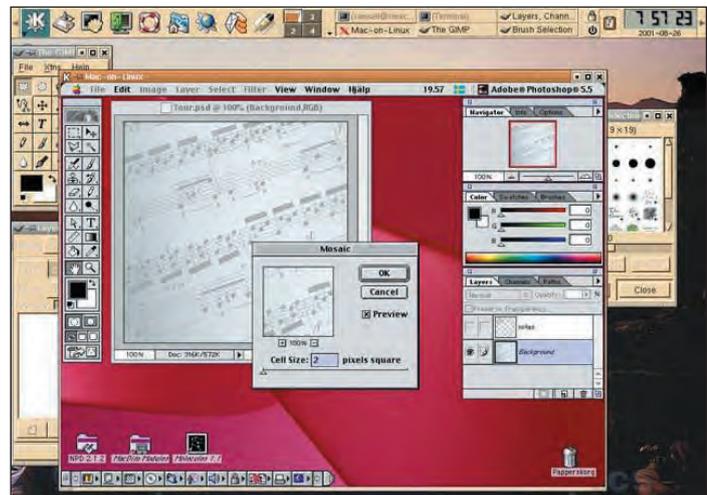
### How to keep up

You may have noticed that some programs come in more than one version. There are two main reasons for this. The first is that the author may have released the latest, experimental, code as a development version while still maintaining the stable release version. In this case, the higher numbered version is likely to have more features, but probably more bugs too.

The other reason, which mainly applies to the DVD, is that the creation of RPM and/or Debian packages for a program is often done by third parties and can lag behind the tarball versions. If you want the latest version of one of these, you'll have to do the `configure;make;make install` thing, or you can have the convenience of an automated install at the expense of the latest features.



Change the way your desktop looks with this selection of themes.



Is it a Mac? No, it's MacOS running on LinuxPPC with the help of *MOL*.

# LinuxFormatCoverdiscCD

## Desktop

Last month's CD and DVD contained a number of themes to use with the main feature, KDE3. This seems a popular inclusion so there are more themes in the Desktop directory this month, and you can expect to see at least a few each month from now on.

Most of the programs on the cover discs come with source code and will run on various processors, although a few will only run on X86 hardware. For a change, this month we give you a program that will only run on a PowerPC equipped computer

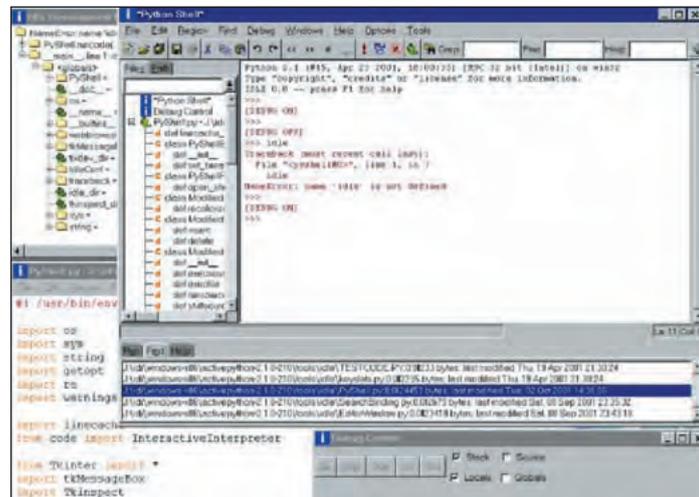
*Mac-on-Linux (MOL)* is a virtual machine that runs MacOS on top of Linux/PPC. There is no CPU emulation

involved since *MOL* runs natively on the PowerPC processor. According to the documentation, all MacOS versions (from 7.5.2 to 9.2.2) are compatible with *MOL*.

## Python

We missed this one when it first came out, mainly because the CD was full of Mandrake 8.2. Python 2.2.1 is an update to the recent Python 2.2 release. It doesn't add new features, it is mainly bug fixes.

By way of compensation for its lateness, we've added some bindings for Qt and KDE as well as *IDE Studio*. This is a Python development environment,



**IDE Studio** is a complete development environment for Python, which is also on the discs this month

## CD CONTENTS AT A GLANCE

### Disc A

#### Internet

<b>Aget</b>	Multi-threaded download accelerator
<b>Anubis</b>	Outgoing mail processor and secure/anonymous SMTP client
<b>Balsa</b>	Gnome e-mail client supporting POP3, IMAP and local mail
<b>cats2procmairc</b>	Generate a .procmairc file from a file of one line rules
<b>DAVIS</b>	Easy and useful language for IRC bots
<b>DilloWebBrowser</b>	Very fast, extremely small Web browser
<b>DownloaderForX</b>	Downloads via both HTTP and FTP with a friendly interface
<b>Evolution</b>	Mailer, calendar, contact manager and communications tool
<b>Gotmail</b>	Non-interactively download email from a Hotmail account
<b>GPG</b>	GNU's tool for secure communication and data storage
<b>jwhois</b>	Configurable Internet Whois client
<b>Mozilla</b>	The second (final?) release candidate for Mozilla 1.0
<b>Netscape</b>	Next generation 7.0 release candidate
<b>Opera</b>	Opera 6.0 has been released, here it is
<b>SecureInstantMessaging</b>	Transparent encryption module for Gaim
<b>Sylpheed</b>	GTK+ based, lightweight, and fast email client
<b>WMBiff</b>	Dockapp to display the status of up to 5 mailboxes

#### Office

<b>ClaraOCR</b>	Optical Character Recognition program
<b>Lazy8Ledger</b>	Very simple yet powerful accounting program
<b>OpenOffice</b>	Version 1.0 of the Open Office suite
<b>OpenOfficeQuickstarter</b>	Start various OpenOffice modules from the KDE System Tray
<b>QHacc</b>	Personal accounting program based on Qt

#### Server

<b>AccountServicesManager</b>	Automate the time consuming tasks of Web hosting companies
<b>Acemmail</b>	Multuser IMAP/POP3 to Web gateway
<b>ApacheCompileHOWTO</b>	Instructions and examples for compiling Apache and modules
<b>BitTorrent</b>	Tool for copying files from one machine to another
<b>Calamaris</b>	Parses the logfiles of a wide variety of Web proxy servers
<b>Camngr</b>	Manages a collection of web cameras
<b>DownloadPHP</b>	Count file downloads and block downloads from another site
<b>GiftWeb</b>	Online gift registry system
<b>INN</b>	Flexible and configurable Usenet/netnews news server
<b>MailingListSynchronizer</b>	Synchronises two or more mailing lists
<b>MHonarc</b>	HTML mail archiving with index, mail thread linking, etc.
<b>Mikrop</b>	Helps mail servers to scan incoming email
<b>MyHeadlines</b>	Adds syndicated headlines to any PHP/MySQL-based website
<b>PHP</b>	Scripting language especially suited for Web development
<b>PostgreSQL</b>	Robust, next-generation, Object-Relational DBMS (ORDBMS)

**ProFTPD**  
**ProMA**  
**SmokePing**  
**WebcamServer**

Proven, high-performance, scalable FTP server  
Web interface for managing users in a ProFTPD/MySQL setup  
Network latency monitor  
Host a video4Linux supported Webcam

### Disc 2 Magazine

<b>Apache</b>	Apache, 1.3.24 and 2.0.36 beta, for the Apache tutorial
<b>Astronomy</b>	Files for this month's Astronomy feature
<b>CGI</b>	The scripts from this month's CGI tutorial
<b>Emulators</b>	All the files mentioned in this month's Emulation article.
<b>HotPicks</b>	All the programs covered in this month's HotPicks section.
<b>Java</b>	The files to go with this month's Java tutorial
<b>Kylix</b>	Example files from the Kylix tutorials
<b>PatchingTheKernel</b>	Everything needed to follow the kernel patching tutorial
<b>Perl</b>	Example scripts from this month's Perl tutorial

### Desktop

<b>3dDesktop</b>	Switches virtual desktops in a seamless 3-dimensional way
<b>AstroCamGtk</b>	Controls Webcams via stepmotor
<b>Biblos</b>	Powerful indexing tool for CDs and local directories
<b>BubblingLoadMonitorApplet</b>	System CPU and memory load monitor for the GNOME panel
<b>Clok</b>	A new way to view time
<b>Cloxten</b>	Digital alarm clock and counter
<b>DOSemu</b>	Run many DOS programs in a Linux shell
<b>Dox</b>	Extensible browser for manpages and HTML documentation
<b>EasyAccessKeyboard</b>	Support for Logitech Easy Access and Internet Keyboards
<b>EasyTar</b>	Simplifies the use of the tar command
<b>FastCD</b>	Expands each "step" of the path entered as a wildcard
<b>FBmodes</b>	X11 display mode generator
<b>fsort</b>	Fast sort program for fixed-length records
<b>GCompris</b>	Complete educational suite for children 3 to 10 years old
<b>GnoChive</b>	GNOME frontend for all common archivers
<b>GNUAwk</b>	Programming language for data-reformatting
<b>GView</b>	Graphical interface for Ghostscript
<b>Gtk-Agenda</b>	Save and share your phone and e-mail listings
<b>IceDock</b>	IceDock gives dockapp support to IceWM and a few other WMs
<b>IceWM</b>	Window manager designed for speed, usability & consistency
<b>jDictionary</b>	Powerful multi-platform dictionary application
<b>MOL</b>	Runs MacOS on top of Linux/PPC
<b>Qt</b>	GUI toolkit for software developers
<b>Themes</b>	A selection of desktop themes and wallpapers
<b>ThreeDFM</b>	3D file manager that uses OpenGL
<b>Voxpak</b>	GUI for playing, editing, renaming voice and fax messages
<b>Wine</b>	Implementation of the Windows 3.x and Win32 APIs
<b>XCounter</b>	X11-based IP traffic monitoring program

based on the standard *Idle* IDE but adding various enhancements.

## Virus repellent

Linux has a reputation for being all but immune to virus attack. The "all but" part means you shouldn't assume you're safe, it doesn't take long to virus check your system to keep it clean.

*Clam AntiVirus* is a command line virus checker, using the virus database from OpenAntivirus.org, that is designed to be run as a cron task. *Clamav* is also able to check files within tar, gzip and zip archives, although it can't handle bzip2 compressed archives at the moment.

## Omnis studio

We give you the chance to try out *Omnis Studio 3.0.1*, the flexible RAD which we recently reviewed (version 3.1, *LXF26*). *Omnis Studio* boosts developer productivity by enabling them to build database-driven apps from a wide range of pre-built components, easily linked together with its high level fourth generation scripting language.

*Omnis Studio's* flexible object model and rich range of pre-built classes could win anyone over to the RAD way of doing things.

The disc contains the installer, manuals and installation instructions (in the file *install.txt*). You will need a promotional code to register, which

## Source RPMs

### Convenient code

Some of the programs on our cover discs are supplied as *.src.rpm* files. You won't be able to install these with the usual `rpm -ihv packagename` as they contain source code rather than the programs themselves. If you have trouble installing from the binary RPM (usually called something like `package-version-i386.rpm`) you may have more luck with the source RPM. This is usually the case if the binary RPM was created for a different

distribution from the one you use. To install from one of these, you must first type

```
rpm --rebuild packagename
```

This will configure and compile the package, much like when you do a `./configure;;make` with a tarball source. Towards the end of the output you will see one or more lines starting

```
Wrote: /usr/src/...
```

These are the binary packages that you have just created, they should be installed in the normal way.

you can get by registering at [www.omnis.net/linuxformat](http://www.omnis.net/linuxformat) and giving the magazine code: LFO03.

The website will also give you details of the upgrade offer to *Studio 3.2*, which is due out on 10<sup>th</sup> July. [LXF](#)

## Development

<b>Dialogos</b>	Automate network communication via a scripting language
<b>FLTK</b>	C++ graphical user interface toolkit for X11 and others
<b>FOX</b>	Popular C++-based toolkit for GUI development
<b>FreePascalCompiler</b>	Turbo Pascal 7.0 and Delphi-compatible 32bit Compiler
<b>Gambas</b>	Graphical development environment based on Basic
<b>GtkDatabox</b>	Display large amounts of numerical data graphically
<b>gtkmm</b>	C++ interface for the popular GUI library GTK+
<b>IDEStudio</b>	Enhanced version of the standard Python Idle IDE
<b>Image-Size</b>	Fetches image dimensions and type from graphics files
<b>JavaCheckstyleTool</b>	Tool to help write Java code to a coding standard
<b>Jregex</b>	Regular expression library for Java
<b>Libsigc++</b>	Callback system for abstract interfaces & general programs
<b>Oberon2</b>	Simple and safe object-oriented language
<b>Omnis Studio</b>	RAD trial - and upgrade offer
<b>PatchUtils</b>	Collection of tools for manipulating patch files
<b>ProgrammerUtilities</b>	Several programs to save the time of developers
<b>pychart</b>	Creates high-quality Postscript or PDF scientific charts
<b>PyKDE</b>	Comprehensive set of Python bindings for KDE 2.x and 3.x
<b>PyQt</b>	Comprehensive set of Python bindings for the Qt GUI toolkit
<b>Python</b>	Bug fix update to Python 2.2
<b>Pyui</b>	User interface library written entirely in Python
<b>ReginaRexxInterpreter</b>	Implementation of the ANSI Standard for the Rexx language
<b>Scriptix</b>	Lightweight scripting language

## Games

<b>Arienne</b>	Massively multiplayer Internet role playing game
<b>Cosmosmash</b>	Destroy rocks before they hit the ground
<b>CrimsonFields</b>	Tactical war game in the tradition of Battle Isle
<b>Danican</b>	Play international draughts
<b>Falsoyd</b>	Entirely plotless game
<b>FreeCellJavaApplet</b>	This is a Java applet of the popular FreeCell card
<b>Gewels</b>	Gnome version of Jewels with multiplayer deadmatch
<b>GL-117</b>	OpenGL- and SDL-based action flight simulator
<b>IceBreaker</b>	Action-puzzle game, capture penguins to ship to Finland
<b>KOPS</b>	Two-dimensional gravitational shoot'yer'friend game
<b>MD3View</b>	OpenGL program to produce skins for Quake 3
<b>nInvaders</b>	nInvaders is a Space Invaders clone for ncurses.
<b>PCSX</b>	PCSX is a working Playstation emulator for Linux.
<b>PizzaBusiness</b>	Pizza Business is a console-based game written in C++.
<b>Stoned</b>	Simple but fully functional curling simulation
<b>Trophy</b>	2D (top-view) action car racing game
<b>UglyGameSearchEngine</b>	Searches the network for game servers
<b>XWelltris</b>	2.5D Tetris-like game

## Graphics

<b>biew</b>	Advanced file viewer with a built-in editor
<b>Giram</b>	Multi-purpose 3D modeller written with GTK+ 1.2
<b>LookAtTheStars</b>	Extended image viewer and manipulator

<b>MjpegTools</b>	Video capture, basic editing, playback, and compression
<b>NewVideoRecorder</b>	High quality video capture toolkit for Linux
<b>Ogle-IR</b>	Infrared Control client for the Ogle DVD player
<b>RayTraccio</b>	Ray tracer handling Constructive Solid Geometry
<b>Scrot</b>	Simple commandline screen capture utility
<b>Slmg</b>	Astronomical image processing software
<b>Terraform</b>	Create and transform fractal terrains
<b>WhiteDune</b>	The standard for displaying 3D data over the web

## Sound

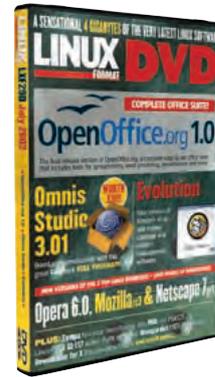
<b>Bonk-XMMS</b>	bonk-xmms is an XMMS input plugin to play bonk files.
<b>BristolSynthesiserEmulator</b>	Synthesizer emulation package
<b>CheeseTracker</b>	Create module music
<b>clxmms</b>	Client of xmms, working in command-line or batch mode
<b>EasyTAG</b>	View and edit tags for MP3, MP2, FLAC and OGG files
<b>Ermixer</b>	Sophisticated OSS mixer with a lot of useful features
<b>Esound</b>	Mixes multiple digitized audio streams and samples
<b>Fische</b>	Standalone sound visualisation
<b>FreeAudioGalaxySatellite</b>	AudioGalaxy client for Linux and BSD
<b>GNUsound</b>	Sound editor supporting multiple tracks
<b>Juked</b>	Turns a server into an MP3 jukebox
<b>Mp3Kult</b>	Organise your MP3/Ogg collection in a MySQL database
<b>OPENDj</b>	Internet radio station software
<b>PlayBonk</b>	playbonk is a spartan ncurses-based bonk player.
<b>Rip-utils</b>	Comprehensive toolkit for managing audio archives
<b>SoundMonitor</b>	GNOME applet to display the output of the esound daemon
<b>TheArtOfNoise</b>	Small but fancy sound editor with plugin support
<b>XMMS-mad</b>	XMMS input plugin that uses libmad to decode MPEG
<b>XVolume</b>	Simple Xlib volume control
<b>YellowStar</b>	Yellow Star is a Communist skin for XMMS.

## System

<b>ClamAntivirus</b>	Powerful anti-virus scanner
<b>Gromit</b>	Automated system configuration tool for multiple machines
<b>GRUB</b>	Multiboot loader
<b>HDparm</b>	View and manipulate various IDE drive and driver parameters
<b>HDup</b>	Back up a filesystem with encryption and compression
<b>HyperactiveDirectoryAdmin</b>	Graphical LDAP tool for administering Users, Groups, Hosts
<b>OpenSSH</b>	Full implementation of the SSH1 protocol
<b>PICA</b>	The Perl Installation and Configuration Agent
<b>Samba</b>	Share printers and files with Linux and Windows machines
<b>Taper</b>	Tape backup and restore program
<b>TranslucencyKernelModule</b>	Kernel module that virtually merges two directories
<b>Watchfolder</b>	Watches specified folders for incoming files
<b>WatchMan</b>	Kills processes that may become a problem
<b>YetAnotherLinuxFAQ</b>	Comprehensive FAQ about the Linux Operating System

# Coverdisc

**Neil Bothwick** is your guide through the wonders of this month's jam-packed *Linux Format DVD*. Hold onto your hats, here we go...



**B**efore you start looking at what's on the DVD, make sure you've checked the CD pages first, as everything on the CDs is also on the DVD. The DVD also allows us to bring you some of the software in different package formats, or to give you additional data (like the star catalogues this month)..

Now let's have a brief look at a few of the items that are only on the DVD. The usual disclaimer applies, this is only a small selection of what is available to you. so make sure you check the full DVD listing on the following page, or put the disc in the drive and load up the index.html page.



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

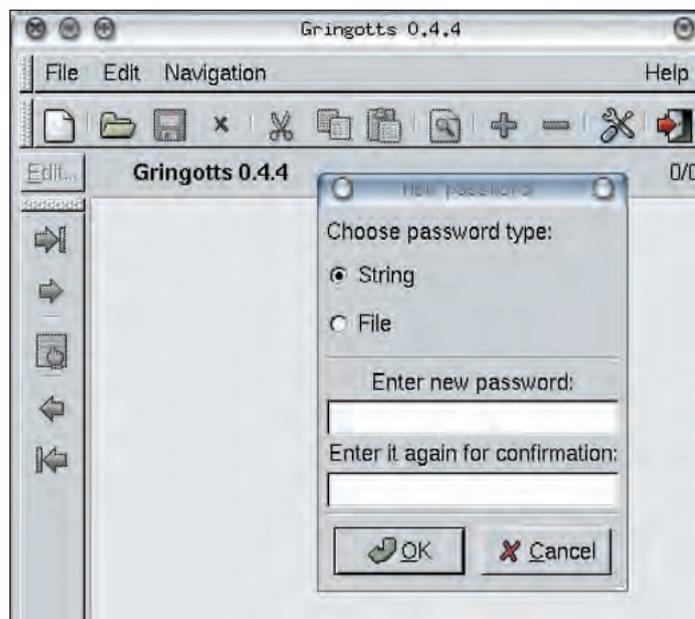
## Important notice

**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.**

## Desktop/Gringotts

We seem to be swamped with passwords nowadays. More and more web sites want you to login, every piece of plastic has a PIN, then there are the login details needed for various online banking sites, and so on. Not only is jotting them on the back of a cigarette packet insecure, there isn't room any more.

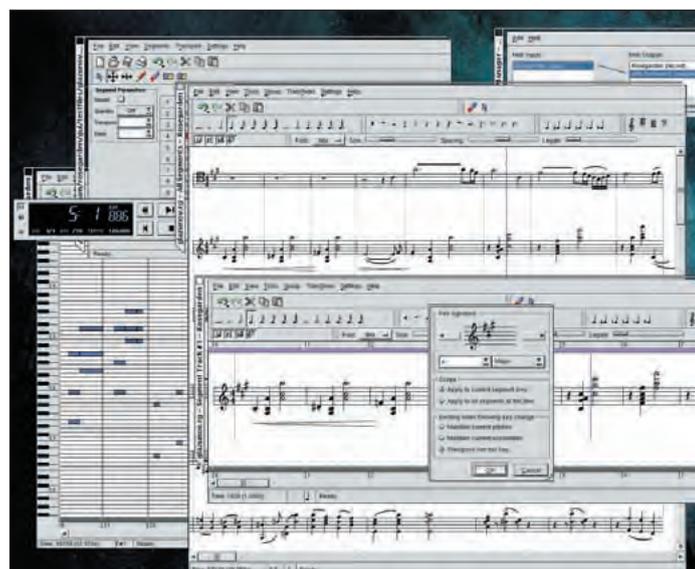
*Gringotts* is a solution to this. It stores all of your secure information in one place, protected by a single passphrase. There are several dependencies, you'll need to install those first. Now you just need to make sure you don't forget the passphrase that unlocks all the others. If you have a Sharp Zaurus, *Keyring* performs a similar function on your PDA.



An electronic fag packet? Store all your important information in one secure place with *Gringotts* or *Keyring*.

## Sound/RoseGarden

*Rose Garden* is a MIDI and audio sequencer using KDE. You don't need to be running KDE to use it, as long as you have the KDE libraries installed. This version adds support for KDE3. In addition to the normal features you would expect to find in a sequencer, *Rose Garden* also provides a musical notation editor. If you have a talent for creating music, or even if you only think you have, this program is worth a look.



Create your own music with the *Rose Garden* sequencer and editor.

## Internet/OutlookTo

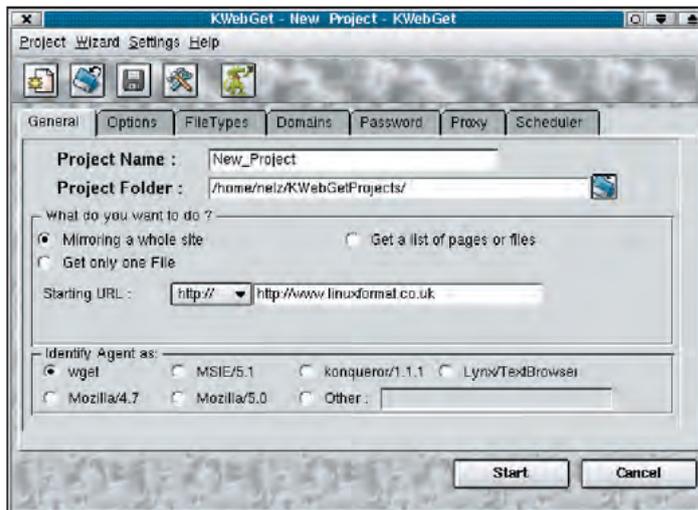
Like it or not (we suspect the latter), a lot of people use *Outlook Express* for email. Once they have a significant amount of mail stored in its own format, they are loath to change mail clients and lose access to it. *OutlookToUnixMailbox* contains tools that convert mail folders saved from *Outlook* as text or CSV into standard mbox files, which can then be read by just about any Linux email client.

## Internet/KWebGet

*Wget* is a powerful way to download individual files or complete directories or sites, using FTP or HTTP. But the

array of command line options can be bewildering to the occasional user.

*KWebGet* is a KDE front end to *wget*, replacing the myriad of command line options with buttons



Grab files, directories or complete web sites with only a few mouse clicks, with *KWebGet* – which works with ftp, as well as webpages.

and checkboxes on several tabbed pages. You can create projects for particular tasks, saving all the settings you use for mirroring a particular site, including passwords. This can save a lot of time if you want to regularly grab files or pages from the same places. Despite the name *KWebGet* can download from FTP sites as well as web sites.

### System/Usermin

We have featured *Webmin* on the coverdiscs a few times. It is a powerful yet easy to use way of administering a Linux system through a web browser. *Usermin* is from the same people and, as the name suggests, is for users. While *Webmin* is mainly concerned with overall system administration, most of its tasks requiring root privileges, *Usermin* allows individual users to easily perform tasks like reading mail, setting up SSH or configuring mail forwarding.

### Graphics/CamStream

*CamStream* is a tool for working with webcams or TV grabber cards. It is fully GUI operated and supports multiple devices, each with a different configuration, so it can handle the output from a camera and a TV card at the same time. Although there are more planned, the features currently implemented include saving in various formats, adjusting image size, brightness and contrast, time lapse recording, adding the date and time to each image and automated uploading to a web server.

### Desktop/MisterHouse

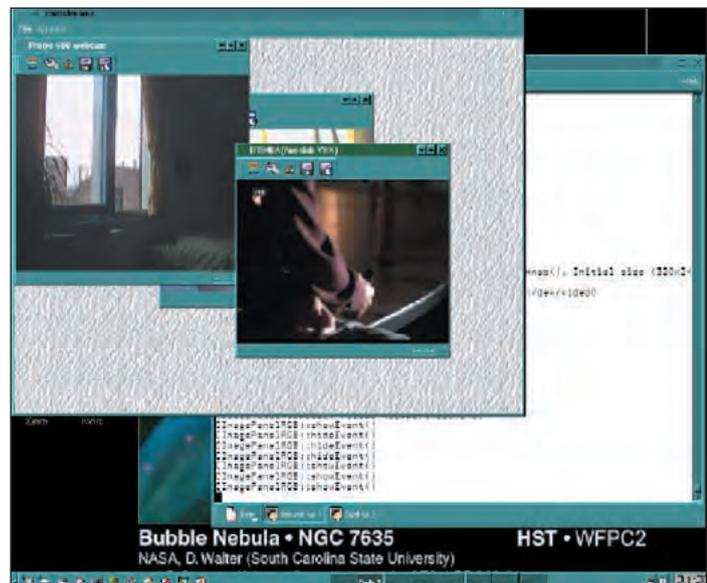
You know the old problem, there are leaves on the line, the train is going to be late and you won't be home in time for that episode of *Coronation Street* you've so been looking forward to. You could phone home and ask someone to tape it, but what if no one else is home? *Mister House* is a home automation system that could deal with this and some serious situations too. It works with standard X10 control devices and can accept commands through the serial port, a web interface or through voice recognition (using *Via Voice*). *Mister House* can also execute events according to a schedule, and all of this can be controlled via the web. OK, so maybe you won't want to use it to tape soap operas, but there are many useful or interesting applications for this sort of automation, see some of the examples in the archive.

### Server/URLEncryptor

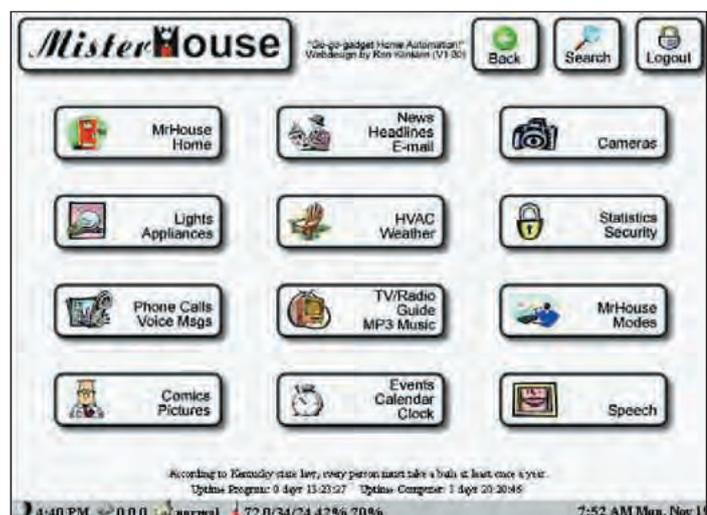
Do you need to keep a list of email addresses on a webpage? It could be a list of contacts or you could be archiving mailing list postings with something like *MHonarc*. If so, you are putting the owners of those addresses in risk of receiving more spam. Spammers use programs that trawl web pages for e-mail addresses. *Otlam* is one way of protecting these addresses while still leaving them available for legitimate use. It scans an HTML document and replaces all mailto: links with an encrypted version, and inserts the



Make life easier for you users, and hopefully yourself, with *Usermin*. The user-space answer to *Webmin*.



Whether you want a security camera or to share images of your cat sleeping with the rest of the world, *CamStream* can help.



Let your computer control your house, but don't leave it in charge of the pod bay doors, Dave.

## LinuxFormatCoverdiscDVD

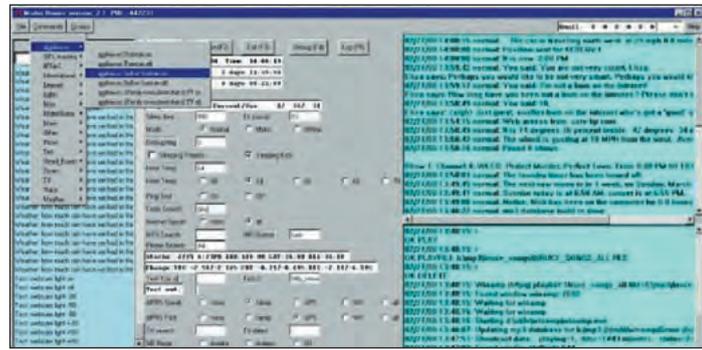
necessary JavaScript code to decrypt them. A visitor using a JavaScript enabled browser will not notice any difference, but the email harvesters will only find garbage. *OtliaM* takes input from stdin, so it can easily be used to encrypt the output of mailing list archivers like *MHonarc*

## Development/SkinLF

Nowadays almost every type of program is skinnable. People like

to be able to tinker with the look of their desktops and the programs that run on them.

As a coder, you may prefer to devote your time to the internal working of your software and adding useful features. *SkinLF* (*Skin Look and Feel*) keeps everyone happy. It is a skin system for Java programs. Using this you can easily add theme capabilities to your software while still leaving time to concentrate on the functionality. [LXF](#)



There's more to *MisterHouse* than remote VCR access.

## DVD CONTENTS AT A GLANCE

## Desktop

<b>Advance</b>	Cross-platform personal information management (PIM) system
<b>BashProgrammableCompletion</b>	Extend bash's completion to achieve complex command lines
<b>CDrecord</b>	The standard cd writing software
<b>FlightGear</b>	Multi-platform, cooperative flight sim development project
<b>FreeGuideTVGuide</b>	Grabs and views TV listings from the Internet
<b>Freetype</b>	High-quality and portable font engine
<b>GPRename</b>	GUI batch file renamer based on Gtk-Perl
<b>Gringotts</b>	Secure storage for passwords, account numbers, PINs, etc.
<b>KAM2</b>	Answering machine for KDE2
<b>Kcd</b>	Directory change utility similar to NCD
<b>KDirStat</b>	Sums up disk usage for direcorey trees
<b>Keyring</b>	Store passwords and secure information on a Zaurus
<b>KSendFax</b>	Interface to fax packages like HylaFax and E-Fax
<b>KTimeclock</b>	Task based timeclock for KDE
<b>MisterHouse</b>	Home automation program written in Perl
<b>NumericalChameleon</b>	Converts numbers with extremely high precision
<b>PerlBox</b>	Desktop manager for UNIX-type systems written in Perl Tk
<b>SimpleBudget</b>	Simplified way to create/manage a budget
<b>WeebleFileManager</b>	PHP file manager that operates through an FTP connection
<b>WMmenu</b>	Dock application to provide a button bar for Windowmaker

## Development

<b>ApacheAnt</b>	Java based build tool, similar to make
<b>COGDevelopment</b>	Software suite for the creation of video games
<b>DirectFB</b>	Thin library that provides hardware graphics acceleration
<b>Frost</b>	Compiler wrapper to use functions with virtual arguments
<b>JSwatJavaDebugger</b>	Graphical, stand-alone Java debugger
<b>LibWMF</b>	Convert WMF (Windows Metafile) format to something useful
<b>OpenCOBOL</b>	Translates COBOL programs to C code to compile with GCC
<b>Python-SIP</b>	Generate C++ interface code for Python
<b>SilverBaloon</b>	OpenGL 3D engine
<b>SkinLF</b>	Provides a skin system for Java programs
<b>SourceHighlight</b>	Produces a syntax highlighted version of a source file

## Distros

<b>Valhalla</b>	ISO images of the Valhalla 7.3 CDs
-----------------	------------------------------------

## Games

<b>Dama</b>	Play Turkish draughts
<b>Groundhog</b>	A small logic game using GTK
<b>Kbilliards</b>	Pool game using Qt/KDE
<b>PackAttack</b>	Move falling boxes into a row to get the best high score
<b>TileWorld</b>	Emulator of "Chip's Challenge"
<b>WeatherMaker</b>	Generate weather reports for role-playing games

## Graphics

<b>CamStream</b>	Streaming video from multiple video sources
<b>Jpegpixi</b>	Interpolates pixels in JFIF images

## LightwaveObjectReader

LWO v2 (Lightwave 6 and above) file reader

## MPlayer

Movie/anim player that supports a wide range of formats

## Xine

Video player for MPEG-1/2, DVD, VCD, SVCD and AVI

## Internet

## CronosII

Powerful GNOME mail client

## Curl

Transfer files using URL syntax

## ElysiumDownloadManager

Extensible download manager for the GNOME desktop

## Everybuddy

Universal Instant Messaging client

## Getleft

Given a URL, Getleft will try to download all links

## KWebGet

KDE based frontend to wget

## OutlookToUnixMailbox

Convert mails saved from Outlook into standard mbox files

## SpamAssassin

Mail filter that uses a wide range of heuristic tests

## Office

## daVinciPresenter

Innovative solution for graph visualization

## phpIntranet

PHP/SQL intranet application with many features

## SOT-Office

Office suite, fully compatible with other major office apps

## Server

## AutomaticRankingPromoter

Search engine Web site promoter

## Blahz-DNS

Blahz-DNS is PHP/MySQL based DNS (BIND 9) administration

## CGI-Application

Create sophisticated, reusable web-based applications

## Dnsmasq

Lightweight DNS forwarder designed for small networks

## DocServ

Indexes documentation in a database

## LinkChecker

Check HTML documents for broken links

## mcNews

Allows visitors to publish News and Comments on your site

## OtliaMURLEncryptor

Encrypts mailto tags in HTML files to stop spam harvesters

## PHPFormMail

PHP script conversion of the original FormMail.pl

## phpMyAdmin

Handle the administration of MySQL over the WWW

## Postfix

Alternative to the widely-used Sendmail program

## Sound

## AlsaPlayer

PCM player for the ALSA sound system

## AlsaPlayer-Python

Write custom interfaces for AlsaPlayer using Python

## ConsoleMixer

No-fuss soundcard mixer application

## Rosegarden

Integrated MIDI sequencer and musical notation editor

## Vorbis

Graphical front-end for the Ogg Vorbis tools

## Xsox

Graphical front-end for sox

## System

## FPauto

Adds automatic Web updates and automatic scans for f-prot

## InstallIDB

Records installs and package information

## Mwvlan

Driver for the WaveLAN/IEEE wireless network card

## SecPanel

GUI for managing and running SSH and SCP connections

## SyslogNG

Syslogd replacement

## Usermin

Simplified version of Webmin for use by normal users

## UtilLinux

Suite of essential utilities for any Linux system

## X86info

x86info probes the registers to find out about your CPU

# Essential disc info

Read this important information before you use your *Linux Format* coverdisc – CD or DVD. We've collated some helpful info to help you get the most from these jewels of data!

## Finding the essentials

### Missing something?

As many of the programs on our discs are the very latest releases, they are often built on the very latest libraries and may depend on other packages your current Linux setup does not contain. We try to provide you with as many of these important supporting files and libraries as possible, though obviously we don't have space to include absolutely everything.

In many cases the latest libraries and

other packages you might need will be included in the "essentials" folder on the disc, so if you are missing dependencies, this is the first place to look.

### Package formats

Wherever possible, we try to include as many different types of package for an installation as possible, whether that be distribution specific RPMs, debs or whatever. Please bear in mind that we can only do this where space permits and when the packages are available.

We will, apart from exceptional or legally restricted situations, include the source files for any package, so that you can build it yourself.

### Documentation

These pages provide helpful information on how to install and use some of the packages on the CD. Please note that many of the applications come with their own documentation, and there are additional notes and files in the relevant directories.

## Creating install CDs with cdrecord

The quickest way to burn an ISO image to CD is with *cdrecord*. You need to be root to do this. First find the address of your CD-writer with

```
cdrecord -scanbus
```

This will show the devices connected to your system. The SCSI address of each device is the three numbers in the leftmost column, say 0,3,0. Now you can burn a CD with

```
cdrecord dev=0,3,0 -v
/path/to/image.iso
```

You can simplify the command by saving some default settings in `/etc/default/cdrecord`. Add a line for each CD writer on your system (usually one) like this

```
Plextor= 0,3,0 12 16M
```

The first item is a label, after the SCSI address you put the speed and the buffer size to use. You can now replace the SCSI address in the command line with the label, but it gets even easier if you add

```
CDR_DEVICE=Plextor
```

Now you can burn an ISO image to disc with

```
cdrecord -v/path/to/image.iso
```

If you really don't want to use the command line, *gcombust* will do the job for you. Start it as root, select the "Burn" tab and the "ISO 9660 Image" gadget near the top of the window. Put the path to the image file in the gadget and press "Combust!". Now put on the kettle while the CD is created for you.

## Other OS?

You do not have to use Linux to burn the ISO to a disc. All the Linux-specific bits are already built into the image file. Programs like *cdrecord* simply dump it to the disk. If you don't have a CD-writer, find someone who does have one, and a DVD drive, and use the CD burning software on their computer. It can be Windows, MacOS, AmigaOS whatever.

## No CD burner?

What if you have no CD writer? Do you know someone else with one? You don't have to use Linux to burn the CDs, any operating system that can run a CD-writer will do the job (see above).

With some distributions it is also possible to mount the images and do a network install, or even a local install from another disk partition. The methods often vary between distributions, so check on the distro vendors website for more info.

## What are all these files?

If you are new to Linux, you may find the profusion of different files and extensions confusing. As we try to give as many packages as possible for compatibility, there will often be two or three files in a directory covering different types of Linux, different architectures and usually source and binary versions – so which do you install? They can be identified by their filenames, and usually just by the file extensions.

Someap-1.01.i386.rpm – This is probably a binary rpm, designed to run on x86 systems.

Someap-1.01.i386.deb – The same, but a debian package.

Someap-1.01.tar.gz – This is usually source code.

Someap-1.01.tgz – Same as the above, tgz is abbreviated form of tar.gz

Someap-1.01.tar.bz2 – Same, but uses bzip2 compression instead of zip

Someap-1.01.src.rpm – This is also source code, but supplied as an rpm to make it easier to install

Someap-1.01.i386.RH7.RPM – A binary, x86 RPM designed specifically for Red Hat Linux

Someap-1.01.ppc.Suse7.rpm – A binary RPM designed specifically for SuSE7.x PPC Linux.

Someap-devel-1.01.i386.rpm – A development version.

## Installing from tarballs

A tar ball is a two stage archive. First the files are archived into a single file with *tar* and then compressed with *Gzip* or *Bzip2*. To unpack, *cd* to the directory you want to unpack it, usually your home directory and type one of the following two lines:

```
tar xzvf /mnt/cdrom/Desktop/progname-2.1.0.tgz
tar xjvf /mnt/cdrom/Desktop/progname-2.1.0.tar.bz2
```

Use the first for Gzipped files, those ending in `.tar.gz` or `.tgz`, and the second for Bzipped files, ending in `.tar.bz2` or `.tbz2`. Naturally, you change the paths to suit the location and name of the archive, and replace `/mnt/cdrom` with whatever is applicable to your system (e.g. `/cdrom`). This normally unpacks the archive into a directory of the same name, enter that directory with:

```
cd progname-2.1.0
```

To compile and install the software, type the following three commands:

```
./configure
make
su -c "make install"
```

The last line will prompt you for the root password, as this stage must be run as root. If you are already logged in as root, just type **make install**. This will give you a default installation. If you want to change any aspect of the install, type **./configure --help** to see the options available. For example, you are usually able to change the default location with the **PREFIX** argument. When you have finished installing, you may remove the source files with:

```
cd ..
rm -fr progname-2.1.0
```

You should also log out as root, before you do anything you may later regret.

## Defective discs

In the unlikely event of your CD/DVD being physically damaged we'll send you a new, working version within 28 days. Send your defective disc – complete with your name, address, and a description of the fault – to: **Linux Format, Future Publishing Disc Department, 3B Athena Avenue, Elgin Industrial Estate, Swindon, SN2 8HF**

# Help wanted

**Free software doesn't just arrive shrink-wrapped and ready on the shelves of some out-of-town retail giant. It's a constant collaborative effort from a community that's happy to have you involved.**

**A**lthough we finish this month's column with education – after diversions through the graphical display side of \*nix – we start it in pure, unadulterated hackerdom, with the openMosix project

You may have read about Mosix in this month's Kernel Patching feature. Mosix allows processes to migrate to other machines on the network – distributing the load for CPU-intensive tasks such as compiling and heavy number-crunching.

When new Mosix releases became proprietary in late 2001, Moshe Bar took up the GPL flag with openMosix. Here he writes on the project, and the help you can contribute to it:

"openMosix is a set of kernel patches and user-space programs to

implement a Single System Image cluster (SSI). An SSI cluster uses the various commodity type PCs connected over standard network links (or better) to create the illusion of one single giant computer with many CPUs.

Next to the SSI functionality we also provide a clustering file system and management tools. openMosix sports already over 1500 installed clusters world-wide with an average of 8000 downloads monthly.

We are always looking for kernel programmers and documentation writers, as well as RPM packagers."

For more details look on the website, or check out the code and get hacking!

<http://openmosix.sourceforge.net/>  
<http://sourceforge.net/projects/openmosix>



## OpenMotif Unix toolkit

**If you are still looking for a** hacking challenge then take a look at *OpenMotif*. Those with boxed sets of Mandrake 8.2 already have the 2.2 release, others can download it from [www.motifzone.org/download](http://www.motifzone.org/download)

*OpenMotif* was started when the *Motif* code was released two years ago. *Motif* was the dominant GUI toolkit on Unix platforms before the growth of Qt and GTK – it is still widely used thanks to its familiarity and cross platform

nature – have a look at some of the products that LXF has reviewed recently and you'll see what we mean.

As the team gear up to prepare for *OpenMotif* 2.3 (release date: Q3 2002), they are calling for open

participation in its development. 2.4 is planned for Q1 2003, and will bring anti-aliased fonts:

[www.motifzone.net/forum/message.php?msg\\_id=2037](http://www.motifzone.net/forum/message.php?msg_id=2037)

Follow the link to the TODO list.

## Utah-GLX Cross-platform 3D

**GLX is an extension to the X** protocol that allows programs to do 3D-rendering with *OpenGL* on an X server. Before the appearance of the Direct Rendering Infrastructure (DRI), with the release of *XFree86 4.0*, *Utah GLX* provided hardware-accelerated 3D drivers for a range of graphics cards under *XFree 3.3.6*. When *XFree86 4.0* took off, the *Utah GLX* project fell by the way side.

The *Utah GLX* project is now being revived and ported to *XFree86 4.0*.

This may seem redundant, but, despite being outperformed by DRI, there are still several valid reasons to use GLX. They supply drivers for cards currently not supported under DRI – such as the ATI RagePro and S3 Savage3D – and they have an open-source 3D driver for NVidia cards. Importantly, since *Utah GLX* doesn't require the complicated kernel direct rendering manager, it is also more portable.

The Utah GLX developers are looking for volunteers to help port the

existing drivers over to *XFree86 4.0*. Currently, ports of the NVidia and RagePro drivers are in progress, but if you interested in getting the other drivers to work and have access to Matrox G200/G400, Intel i810, SiS 6326 or Savage3D hardware, then they want want to here from you. Help with testing and squashing bugs would also be greatly appreciated.

<http://utah-glx.sourceforge.net/>  
<http://sourceforge.net/projects/utah-glx/>

## Your help wanted

There are thousands of Free Software projects in need of some sort of help and thousands of Linux Format readers who may be interested in assisting your project. However we cannot publicise what we don't know about.

If YOU have a project that's in need of anything from artists and beta-testers to web-designers and, er, something beginning with Z, we want to hear about it. Email us now at [linuxformat@futurenet.co.uk](mailto:linuxformat@futurenet.co.uk) and give us some details of your project, and what sort of help you are looking for. Please include plenty of info!

## Education, education, education

### First steps in keeping it Free

**Recent developments, particularly** in Latin America and parts of Europe, have given fresh heart to the diverse efforts to get children educated with Free Software.

For a variety of reasons governments are examining the use of Free Software in various departments, and particularly across the education system. One reason is concern about the amount of tax "dollars" flowing overseas to monopoly software producers for products that they are not actually buying, merely paying to have on their computers for a short period of time. Cost aside, security concerns and, most heartening of all, a real understanding of the importance of freedom, have lead to administrations examining GNU/Linux for not just server duties, but the desktop and the classroom. It is against this backdrop that various educational projects find themselves needing your help to deliver the goods for Free Software's place in schools and colleges.

You can find many separate educational projects (over a hundred) on the 'net involving Free Software and/or Linux. Until recently many of them have duplicated each other's efforts. Now, with Schoolforge (see box), things are becoming more coordinated, but all the projects need more willing volunteers.

### GNU

The GNU project is at the heart of Free Software development but its education website –

[www.gnu.org/education](http://www.gnu.org/education) – needs some help with updating and developing. Contact Odile Bénassy: [obenassy@free.fr](mailto:obenassy@free.fr)

As well as supporting a free encyclopedia [www.gnu.org/encyclopedia/index.html](http://www.gnu.org/encyclopedia/index.html), they partner with OFSET (see LXF 28) in the development of Free educational software. Run from South America, this project has a global outlook on education.

### AFFS

The Association For Free Software (AFFS) is looking for input and help for its campaign to get GNU/Linux

and Free Software used in UK schools, and for IT education to move away from its recent "Point 'n' drool" tendencies.

[www.affs.org.uk/education](http://www.affs.org.uk/education)

AFFS is building partnerships with other education organisations and looking to increase active membership by recruiting from those who wish to improve education in the UK by the inclusion of the values of community and sharing embodied in Free Software – and also those who resent seeing their taxes flying to overseas software monopolists at the expense of home-grown coding talent!

If you want to get involved – through writing, lobbying, testing, distributing, coding, or in any other way – then please get on the mailing list and make yourself known.

[www.affs.org.uk/contact.html](http://www.affs.org.uk/contact.html)

### Free Europe

Meanwhile the FSFE's education pages and mailing list act as a meeting point for Europe's different projects. If you want to share your knowledge or ideas – or simply find out what's going on – sign up for the mailing list at [http://mail.gnu.org/mailman/listinfo/edu-eu/](http://mail.gnu.org/mailman/listinfo/edu-eu)

### debian-jr

With the imminent release of Woody (Debian 3.0), the debian-jr project will get its long awaited first airing. Using Debian's flexible and powerful *taskel* installation, you will be able to choose packages suitable for computer users aged 2-10, and their administrators.

Even though this is just the beginning for Debian-jr – the project aims at a complete system suitable for children – what has been packaged for Woody will be useful to any parent (and most teachers too).

This first release, with *taskel* selection and a few dozen basic apps, is aimed at just getting Debian-jr out there in the wild and getting feedback and ideas, so that more people (and not just developers) hop on board.



Help is always needed – you don't have to be a Debian developer, join in selecting and documenting packages; test them or help with the website [www.debian.org/devel/debian-jr/jobs](http://www.debian.org/devel/debian-jr/jobs)

### KDE

*KDE edu* is KDE's edutainment project, designed to create free educational software based around KDE, the K Desktop Environment, for children aged 3 to 18 and more.

Currently, there are few educational applications within KDE but from the KDE3.0 release, the KDE Edutainment package is now released with the rest of the official KDE – which should give a big boost to their efforts.

What can you do to help? Well if you can code C++ and know the KDE libs, head over to <http://edu.kde.org/ideas/index.phtml> and look at the suggested projects.

Non-coders should head over to the mailing list and make known their other skills.

### Seul

The SEUL/edu project – instrumental in setting up the schoolforge project (see *Schoolforge* box) – is working on an educational ISO.

"This project will develop an easily installable ISO image file (CDROM) of

## Schoolforge

Schoolforge was set up earlier this year to minimise all the duplication worldwide between different lists promoting Linux and open learning resources into schools. Although, like SEUL/edu, it is not solely focused on Free Software, it is a valuable resource – being made up of various groups around the world with different perspectives on education.

Schoolforge is looking for success stories for its case studies and translators for its press releases. <http://schoolforge.net/>

educational software and supporting material that runs on Linux. It will be distribution-neutral to the greatest extent possible. The applications will be grouped in such a way that users can choose standard selections of software for various ages and courses of study without having to select each package individually, although they'll be able to do so if they desire."

They need help from parents, children, schools and teachers to select and review suitable apps – the more the better. Reviewing the applications doesn't require a lot of deep technical expertise; in fact it's probably better if the reviewers aren't technical experts.

Email Doug Loss: [dloss@seul.org](mailto:dloss@seul.org)  
[www.seul.org/edu/iso.html](http://www.seul.org/edu/iso.html)

# User Groups

Your local Linux User Group needs you! LUGs worldwide are 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.lug.org.uk](http://www.lug.org.uk) or <http://lugwww.counter.li.org/groups.cms>

## 1 Hampshire

URL [www.hants.lug.org.uk](http://www.hants.lug.org.uk)  
Contact Hugo Mills

## 2 Bristol & Bath

URL [www.bristol.lug.org.uk](http://www.bristol.lug.org.uk)

## 3 Scottish

URL [www.scottish.lug.org.uk](http://www.scottish.lug.org.uk)  
Contact Tony Dyer

## 4 Oxford

URL [www.oxford.lug.org.uk](http://www.oxford.lug.org.uk)  
Contact Alasdair G Kergon

## 5 Kent

URL [www.kent.lug.org.uk](http://www.kent.lug.org.uk)  
Contact John Mills

## 6 Brighton

URL [www.brighton.lug.org.uk](http://www.brighton.lug.org.uk)  
Contact Johnathan Swan

## 7 Worcestershire

URL [www.worcs.lug.org.uk](http://www.worcs.lug.org.uk)  
Email [info@thirdeyedevlopment.com](mailto:info@thirdeyedevlopment.com)

**NEW  
DETAILS**

## 8 Northants

URL [www.northants.lug.org.uk](http://www.northants.lug.org.uk)  
Contact Kevin Taylor

## 9 Anglian

URL [www.anglian.lug.org.uk](http://www.anglian.lug.org.uk)  
Contact Martyn Drake

## 10 Milton Keynes

URL [www.mk.lug.org.uk](http://www.mk.lug.org.uk)  
Contact Denny De La Haye

## 11 Doncaster

URL [www.doncaster.lug.org.uk](http://www.doncaster.lug.org.uk)  
Contact Andy Smith

## 12 Moray

URL [www.moray.lug.org.uk](http://www.moray.lug.org.uk)  
Contact Stewart Watson

## 13 West Wales

URL [www.westwales.lug.org.uk](http://www.westwales.lug.org.uk)  
Contact Dan Field

## 14 Wolves

URL [www.wolves.lug.org.uk](http://www.wolves.lug.org.uk)  
Contact Jono Bacon

## 15 Peterborough

URL [www.peterboro.lug.org.uk](http://www.peterboro.lug.org.uk)  
Contact Steve Gallagher

## 16 Edinburgh

URL [www.edinburgh.lug.org.uk](http://www.edinburgh.lug.org.uk)  
Contact Alistair Murray

## 17 Tyneside

URL [www.tyneside.lug.org.uk](http://www.tyneside.lug.org.uk)  
Contact Brian Ronald

## 18 Leicester

URL [www.leicester.lug.org.uk](http://www.leicester.lug.org.uk)  
Contact Clive Jones

## 19 Greater London

URL <http://gl.lug.linux.co.uk/>  
Contact John Southern

## 20 Surrey

URL [www.surrey.lug.org.uk](http://www.surrey.lug.org.uk)  
Contact Jay Bennie

## 21 Cambridge

URL [www.cam-lug.org](http://www.cam-lug.org)

## 22 Devon & Cornwall

URL [www.dclug.org.uk](http://www.dclug.org.uk)  
Contact Simon Waters

## 23 Falkirk

URL [www.falkirk.lug.org.uk](http://www.falkirk.lug.org.uk)

## 24 Manchester

URL [www.manlug.mcc.ac.uk](http://www.manlug.mcc.ac.uk)  
Contact John Heaton, Owen Le Blanc

## 25 Hertfordshire

URL [www.herts.lug.org.uk](http://www.herts.lug.org.uk)  
Contact Nicolas Pike

## 26 West Yorkshire

URL [www.wylug.lug.org.uk](http://www.wylug.lug.org.uk)  
Contact Jim Jackson

## 27 Sheffield

URL [www.sheflug.co.uk](http://www.sheflug.co.uk)  
Contact Richard Ibbotson

## 28 Staffordshire

URL [www.staffslug.org.uk](http://www.staffslug.org.uk)

## 29 North East

URL [www.shofar.uklinux.net/NELUG](http://www.shofar.uklinux.net/NELUG)

## 30 London

URL [www.lonix.org.uk](http://www.lonix.org.uk)

## 31 Thames Valley

URL [www.sclug.org.uk](http://www.sclug.org.uk)

## 32 Liverpool OpenSource

URL [http://linux.liv.ac.uk/\\_liv\\_linux\\_ug/](http://linux.liv.ac.uk/_liv_linux_ug/)  
Contact Simon Hood

## 33 Deal Amiga Club

Email [superhighwayman@hotmail.com](mailto:superhighwayman@hotmail.com)  
Contact John Worthington

## 34 Chesterfield

Email [spirelug@yahoo.co.uk](mailto:spirelug@yahoo.co.uk)  
Contact Robin Needham

## 35 South Derbyshire

URL [www.sderby.lug.org.uk](http://www.sderby.lug.org.uk)  
Contact Dominic Knight

## 36 Belfast (BLUG)

URL [www.belfastlinux.cx](http://www.belfastlinux.cx)  
Contact Ken Guest

## 37 Wiltshire

URL [www.wiltshire.lug.org.uk](http://www.wiltshire.lug.org.uk)  
Contact Jason Rudgard

## 38 South London

URL [www.sl.lug.org.uk](http://www.sl.lug.org.uk)  
Email [ben@ilovephilosophy.com](mailto:ben@ilovephilosophy.com)

**REVISED  
DETAILS**

## 39 Cheshire

URL [www.sc.lug.org.uk](http://www.sc.lug.org.uk)  
Contact Anthony Prime – [enquiry@sc.lug.org.uk](mailto:enquiry@sc.lug.org.uk)

## 40 North Wales

URL [www.northwales.lug.org.uk](http://www.northwales.lug.org.uk)  
Contact Jonathan Cole

## 41 Midlands

URL [www.midlandsLUG.cjb.net](http://www.midlandsLUG.cjb.net) **WARNING: Popup ads**  
Contact Pete Thompson

## 42 Cumbria

URL [www.cumbria.lug.org.uk](http://www.cumbria.lug.org.uk)  
Contact Jamie Dainton

## 43 Dorset

URL [www.dorset.lug.org.uk](http://www.dorset.lug.org.uk)  
Contact John and Mat

**REVISED  
DETAILS**

## 44 Shropshire

URL [www.shropshire.lug.org.uk](http://www.shropshire.lug.org.uk)  
Email [shropshire@lug.org.uk](mailto:shropshire@lug.org.uk)

## 45 South West

URL [www.southwestlug.uklinux.net](http://www.southwestlug.uklinux.net)  
Email [southwest@lug.org.uk](mailto:southwest@lug.org.uk)

## 46 South Wales

URL [www.sw.lug.org.uk](http://www.sw.lug.org.uk)  
Contact Tim Bonnell

## 47 North London

URL <http://www.kemputing.net/lug/anlug-aims.html>

## 48 Malvern

URL [www.malvern.lug.org.uk](http://www.malvern.lug.org.uk)  
Contact Greg Wright

## 49 Huddersfield

URL [www.hud.lug.org.uk](http://www.hud.lug.org.uk)  
Contact Adam Brookes

## 50 Nottingham

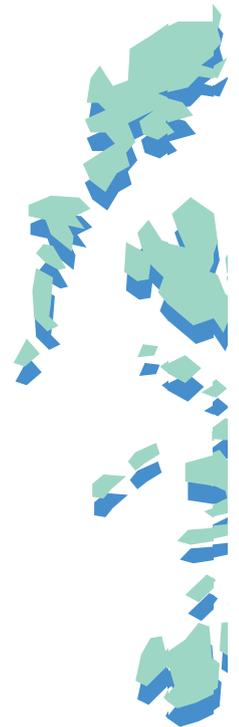
URL [www.nottingham.lug.org.uk](http://www.nottingham.lug.org.uk)  
Contact Godfrey Nix

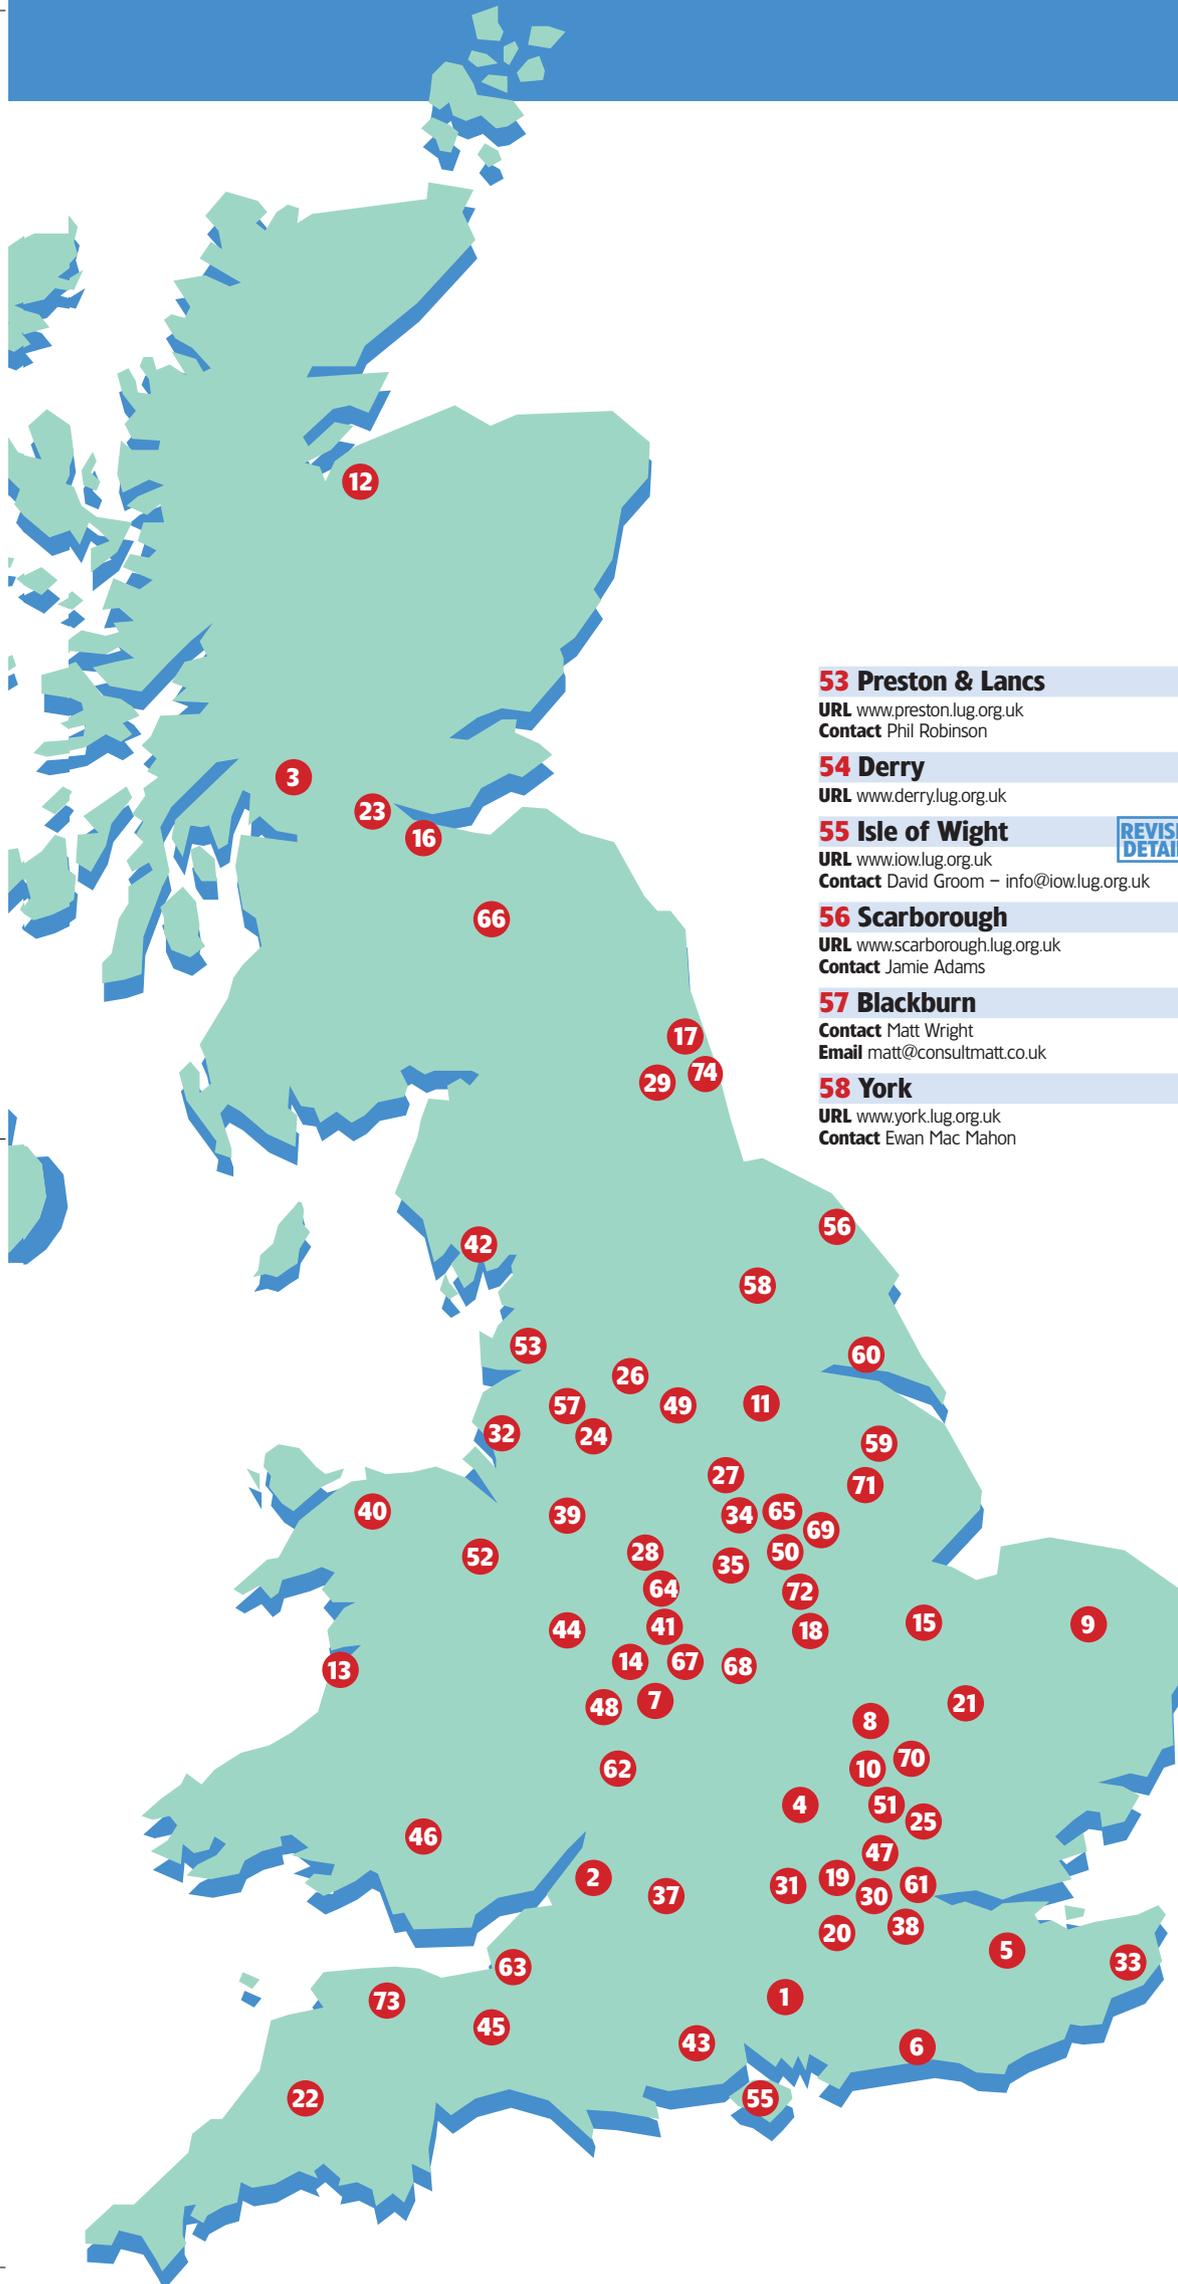
## 51 St Albans & Luton

URL [www.lust.lug.org.uk](http://www.lust.lug.org.uk)  
Contact Michael Culverhouse – [mike@easily.co.uk](mailto:mike@easily.co.uk)

## 52 Wrexham

Contact Paul Kersey-Smith  
Email [paul@pkls.fsnet.co.uk](mailto:paul@pkls.fsnet.co.uk)



**53 Preston & Lancs**

URL [www.preston.lug.org.uk](http://www.preston.lug.org.uk)  
 Contact Phil Robinson

**54 Derry**

URL [www.derry.lug.org.uk](http://www.derry.lug.org.uk)

**55 Isle of Wight**

URL [www.iow.lug.org.uk](http://www.iow.lug.org.uk)  
 Contact David Groom – [info@iow.lug.org.uk](mailto:info@iow.lug.org.uk)

REVISED  
DETAILS**56 Scarborough**

URL [www.scarborough.lug.org.uk](http://www.scarborough.lug.org.uk)  
 Contact Jamie Adams

**57 Blackburn**

Contact Matt Wright  
 Email [matt@consultmatt.co.uk](mailto:matt@consultmatt.co.uk)

**58 York**

URL [www.york.lug.org.uk](http://www.york.lug.org.uk)  
 Contact Ewan Mac Mahon

**59 Lincs**

URL [www.lincs.lug.org.uk](http://www.lincs.lug.org.uk)  
 Contact Chris Lingard

**60 Hull**

URL <http://hulllug.cjb.net>  
 Contact [hulllug@jaxnetwork.co.uk](mailto:hulllug@jaxnetwork.co.uk)

REVISED  
DETAILS**61 East London**

URL [www.eastlondon.lug.org.uk](http://www.eastlondon.lug.org.uk)  
 Contact Jonathan Spriggs

**62 Gloucestershire & Cotswolds**

URL [www.gloucs.lug.org.uk](http://www.gloucs.lug.org.uk)  
 Contact Barrie Haycock

**63 Yeovil College**

URL [www.yc.lug.org.uk](http://www.yc.lug.org.uk)  
 Contact Adam Parker

**64 South Staffordshire**

URL [www.staffs.lug.org.uk](http://www.staffs.lug.org.uk)  
 Contact Oliver Keenan

**65 Mansfield**

URL [www.mansfield.lug.org.uk](http://www.mansfield.lug.org.uk)  
 Contact Brent Vardy

**66 Borders**

URL [www.linux.bordernet.co.uk](http://www.linux.bordernet.co.uk)  
 Contact Welby McRoberts

**67 South Birmingham**

URL [www.sb.lug.org.uk](http://www.sb.lug.org.uk)  
 Contact Tim Williams

**68 Coventry**

Contact Darren Austin  
 Email [info@coventrylug.org.uk](mailto:info@coventrylug.org.uk)

**69 Newark & Lincoln**

URL [www.newlinc.lug.org.uk](http://www.newlinc.lug.org.uk)

**70 Bedfordshire**

URL [www.beds.lug.org.uk](http://www.beds.lug.org.uk)  
 Contact Neil Darlow

**71 Lincoln**

URL [www.lincoln.lug.org.uk](http://www.lincoln.lug.org.uk)  
 Contact Jon Shamash

**72 Loughborough**

URL [www.loughborough.lug.org.uk](http://www.loughborough.lug.org.uk)  
 Contact Martin Hamilton

**73 Exeter University**

Contact Nicholas Murison  
 Email [N.J.Murison@exeter.ac.uk](mailto:N.J.Murison@exeter.ac.uk)

**74 Sunderland**

Contact Thomas Croucher  
 Email [thomas.croucher@sunderland.ac.uk](mailto:thomas.croucher@sunderland.ac.uk)

# LinuxUserGroups

## LUG OF THE MONTH!

### Manchester

**ManLUG, the Manchester Linux User Group**, began to meet in 1994, largely at the instigation of Dr David (Nobby) Clark. We have quite active mailing lists, monthly meetings, and a permanent home, thanks to the University of Manchester. Some of our members are absolute beginners, some networking or kernel gurus, and some at every level in between. Many of us volunteer to help local Linux users to solve their problems.

One founding member, John Heaton, maintains our website – here you can find information about the group, our history, coming events, mailing lists, and archives. There are also many pictures and the names and email addresses of people to contact for further information. Some ManLUG members frequently attend meetings of other Linux user groups.

ManLUG meets on the third

Saturday of every month except December. We usually have two subjects per meeting, and try to make them as different from each other as possible, so that meetings appeal to as many people as possible. There is no fee and no 'official' list of members, though the University security staff may ask people for identification.

[www.manlug.mcc.ac.uk](http://www.manlug.mcc.ac.uk)



## Worldwide Linux User Groups

Free Software users across the globe

### Australia

#### ADELAIDE

**URL** [www.linuxsa.org.au](http://www.linuxsa.org.au)  
**Email** [mtippet@anu.edu.au](mailto:mtippet@anu.edu.au)

#### MELBOURNE, VICTORIA

**URL** [www.luv.asn.au](http://www.luv.asn.au)  
**Contact** [luv-committee@luv.asn.au](mailto:luv-committee@luv.asn.au)

#### PERTH

**URL** <http://plug.linux.org.au/>

### Europe

#### AUVERGNE

**URL** [www.linux-arverne.org](http://www.linux-arverne.org)  
**Email** [Cyril.Hansen@wanadoo.fr](mailto:Cyril.Hansen@wanadoo.fr)

#### COSTA DEL SOL (English speaking)

**URL** [www.fuengirola.lug.org.uk](http://www.fuengirola.lug.org.uk)

#### DENMARK

**AIssund** [www.alslug.dk](http://www.alslug.dk)

**Esbjerg** [www.eslug.dk](http://www.eslug.dk)

**Fyns** [www.flug.dk](http://www.flug.dk)

**Midt-og Vestjylland** [www.mvjlug.dk](http://www.mvjlug.dk)

**Nordjylland** [www.njlug.dk](http://www.njlug.dk)

**Skåne Sjælland** [www.sslug.dk](http://www.sslug.dk)

**Trekantsområdet** [www.tlug.dk](http://www.tlug.dk)

**Vest-fyn** [www.haarby-net.dk/vflug](http://www.haarby-net.dk/vflug)

**Århus** [www.aalug.dk](http://www.aalug.dk)

#### EIRE

**URL** [www.linux.ie](http://www.linux.ie)  
**Email** [root@linux.ie](mailto:root@linux.ie)  
**URL** [www.dilu.org](http://www.dilu.org)  
**Email** [glossary@dilu.org](mailto:glossary@dilu.org)

#### GOTHENBURG

<http://nain.oso.chalmers.se/LUGG/>

**UK** Don't forget the distro-specific lists:

**URL** [www.lug.org.uk/maillist.html](http://www.lug.org.uk/maillist.html)

### India

**URL** [www.linux-india.org](http://www.linux-india.org)  
**Email** [newsmaster@linux-india.org](mailto:newsmaster@linux-india.org)

#### TRIVANDRUM

**URL** [www.river-valley.com/tux](http://www.river-valley.com/tux)  
**Email** [anil@river-valley.com](mailto:anil@river-valley.com)

### Middle East

#### EGYPT

**URL** [www.linux-egypt.org](http://www.linux-egypt.org)  
**Contact** Hesham Bahram

### North America

#### ALASKA

**URL** [www.aklug.org](http://www.aklug.org)  
**Email** [deem@wdm.com](mailto:deem@wdm.com)

#### BATON ROUGE

**URL** [www.br lug.net](http://www.br lug.net)  
**Email** [dpuryear@usa.net](mailto:dpuryear@usa.net)

#### BAY AREA

**URL** [www.balug.org](http://www.balug.org)  
**Email** [afyde@balug.org](mailto:afyde@balug.org)

#### CLARKSVILLE, TN

**URL** [www.clug.org](http://www.clug.org)  
**Email** [tux@clug.org](mailto:tux@clug.org)

#### DENVER

**URL** <http://clue.denver.co.us/>  
**Email:** [president@clue.denver.co.us](mailto:president@clue.denver.co.us)

#### FLORIDA

**URL** [www.flux.org](http://www.flux.org)

#### LOS ANGELES

**URL** [www.lalugs.org](http://www.lalugs.org)  
**Email** [dank@alumni.caltech.edu](mailto:dank@alumni.caltech.edu)

#### NORTH COLORADO

**Email** [nclug@nclug.org](mailto:nclug@nclug.org)  
**Contact** Mat Taggart

#### TAMPA

**URL** [www.suncoastlug.org](http://www.suncoastlug.org)  
**Email** [president@suncoastlug.org](mailto:president@suncoastlug.org)

#### UHACC Normal, IL

**URL** [www.uhacc.org](http://www.uhacc.org)  
**Email** [lug@uhacc.org](mailto:lug@uhacc.org)

#### VIRGINIA TECH

**URL** [www.vtluug.org](http://www.vtluug.org)  
**Email** [nega@vt.edu](mailto:nega@vt.edu)

### South America

#### BUENOS AIRES

**Email** [dcoletti@impost.com.ar](mailto:dcoletti@impost.com.ar)

#### CHILE

**URL** [www.linux-chile.org](http://www.linux-chile.org)

#### MONTEVIDEO

**URL** [www.linux.org.uy](http://www.linux.org.uy)

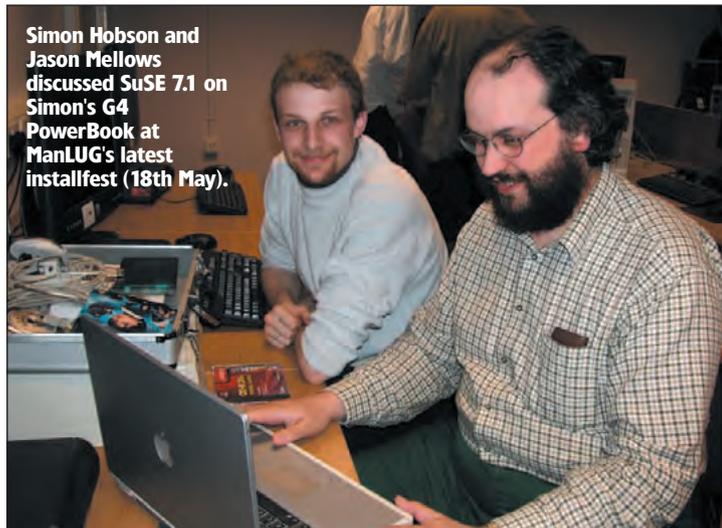
#### PARAGUY/ ASUNCION

**Email** [rolgiati@conexion.com.py](mailto:rolgiati@conexion.com.py)

#### SAO PAULO

**URL** <http://gul.ime.usp.br/>  
**Email** [gul@ime.usp.br](mailto:gul@ime.usp.br)

Simon Hobson and Jason Mellows discussed SuSE 7.1 on Simon's G4 PowerBook at ManLUG's latest installfest (18th May).



## Linux Install Day 2002

"Newcomers, old hats and gurus all mixed in a friendly atmosphere," writes Christopher Birchenhall. "John Heaton installed Red Hat 7.2 from the network and demonstrated the multimedia capacities of his own laptop running SuSE 7.2. Owen LeBlanc helped Vince and Jonathan Wrigley load Mandrake 8.2; Fred Dodd proffered additional advice on firewalls, and assisted Martin Hickley installing Mandrake 8.2 and KDE 3.0 onto Martin's laptop. Owen helped me with SuSE 8.0 on my laptop. Tobias Schiebeck explained his

switch from SuSE to Debian; he had found *YaST* disturbed his manually set configurations, while relying on the command line interface in Debian left him fully in control. Kizito Birabwa outlined his own unified HOWTO on Linux Locales as part of his efforts to introduce Uganda's native language (Luganda) into Linux, through the Translation Project (University of Montreal). Several attendees, including myself, walked away with raffle prizes including free boxed copies of SuSE 7.3 and several O'Reilly books!"

[www.manlug.mcc.ac.uk](http://www.manlug.mcc.ac.uk)

## Linux User Group organisers

If you're not listed here, or we have your details wrong, please contact us at: **LUGS!, Linux Format, 30 Monmouth Street, Bath, BA1 2BW** or email your details to: [linuxformat@futurenet.co.uk](mailto:linuxformat@futurenet.co.uk)

**EDITORIAL****Editor** Nick Veitch [nick.veitch@futurenet.co.uk](mailto:nick.veitch@futurenet.co.uk)**Art Editor** Julian Jefferson [julian.jefferson@futurenet.co.uk](mailto:julian.jefferson@futurenet.co.uk)**Reviews Editor** Richard Drummond [richard.drummond@futurenet.co.uk](mailto:richard.drummond@futurenet.co.uk)**Production Editor** Richard Smedley [richard.smedley@futurenet.co.uk](mailto:richard.smedley@futurenet.co.uk)**Editorial Contributors**

Jono Bacon, Neil Bothwick, Chris Brown, David Cartwright, Paul Cavanagh, Andy Channele, David Coulson, Dave Cross, Hoyt Duff, Simon Goodwin, Steve Heaven, Biagio Lucini, Maurice Kelly, Brian Long, Mike Saunders, Charlie Stross, Tom Wilkinson

**COVER CD PRODUCTION****CD Editor** Neil Bothwick**Senior CD Editor** Kate Hadley**ART CONTRIBUTORS****Photography** Chris Birchenhall, Amanda Thomas, James Wilson, Getty Images**Illustration** Paul Bateman, Chris Winn, Shane Collinge**ADVERTISING SALES**

Matt Stanley 01225 442244

[matt.stanley@futurenet.co.uk](mailto:matt.stanley@futurenet.co.uk)

Emma Mimmack 01225 442244

[emma.mimmack@futurenet.co.uk](mailto:emma.mimmack@futurenet.co.uk)**MARKETING AND PROMOTIONS****Marketing Manager** Fiona Tully**Subscriptions Executive** Nick Skardon**PRODUCTION****Production Co-ordinator** Diane Ross**Group Production Manager** Clare Tovey**Advertising Production Co-ordinator** Jo Crosby**MANAGEMENT****Publisher** Dave Taylor**Group Advertising Manager** Simon Moss**Group Publisher** John Weir**Circulation Director** Sue Hartley**Managing Director** Colin Morrison**DISTRIBUTION AND CIRCULATION****Circulation Manager** Jamie Malley[jamie.malley@futurenet.co.uk](mailto:jamie.malley@futurenet.co.uk)**Distributed by** Seymour Distribution, 86 Newman Street, London W1T 3EX**Tel** +44 (0)207 3968000**Overseas Distribution** by Future Publishing Ltd.**Tel** 01225 442244.**Overseas Licences**Simon Wear [simon.wear@futurenet.co.uk](mailto:simon.wear@futurenet.co.uk)**Tel** +44(0)1225 442244 **Fax** +44 (0)1225 732361**Contact Details**

Linux Format, 30 Monmouth Street, Bath BA1 2BW

**Tel** +44 (0)1225 442244 **Email** [linuxformat@futurenet.co.uk](mailto:linuxformat@futurenet.co.uk)**Subscriptions and Mail Order****Phone** +44 (0)1458 271178. See page 96**Email** [linuxformat.subs@futurenet.co.uk](mailto:linuxformat.subs@futurenet.co.uk)**Copyright** No part of this publication may be reproduced without written permission from our publisher. We assume all letters sent – by email, fax or post – are for publication unless otherwise stated, and reserve the right to edit contributions. All contributions to Linux Format are submitted and accepted on the basis of

non-exclusive worldwide licence to publish or license others to do so unless otherwise agreed in advance in writing recognises all copyrights and trademarks. Where possible, we have acknowledged the copyright holder. Contact us if we haven't credited your copyright and we will always correct any oversight.

**All CD-ROM** demos and reader submissions are supplied to us on the assumption they can be incorporated into a future covermounted CD-ROM, unless expressly stated to the contrary. We cannot be held responsible for mistakes or misprints. Linux Format recognises all copyrights in this issue. Where possible we have acknowledged the copyright holder. Please contact us if we have failed to credit copyright.**Disclaimer** All tips in this magazine are used at your own risk. We accept no liability for any loss of data or damage to your computer, peripherals or software through the use of any tips or advice.**Printed** in the UK by Midway Clark (Holt) and © Future Publishing Ltd 2002**LINUX** is a trademark of Linus Torvalds, GNU/Linux is abbreviated to Linux throughout for brevity.

All other trademarks are the property of their respective owners

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

The Future Network produces carefully targeted specialist magazines and websites for groups of people who share a passion. We aim to satisfy their passion by creating titles that offer superb value for money, trustworthy information, multiple ways to save time and money, and are a pleasure to read or visit. Today we publish more than 80 magazines and over 30 magazine websites and networks from offices in five countries. The company also licenses 32 of its titles resulting in over 60 local editions in a further 23 countries. The Future Network plc is a public company quoted on the London Stock Exchange (symbol: FNET).

**Non-executive Chairman:** Roger Parry**Chief Executive:** Greg Ingham**Chief Operating Officer & Managing Director, UK:** Colin Morrison**Group Finance Director:** John Bowman**Tel** +44 (0)1225 442244[www.thefuturenetwork.plc.uk](http://www.thefuturenetwork.plc.uk)BATH • LONDON • MILAN •  
NEW YORK • PARIS •  
SAN FRANCISCO • WROCLAW

# NEXT MONTH

Issue 30 on sale Thursday 18 July

The ultimate guide to

# OFFICE SOFTWARE

Not ready for the desktop? Don't you believe it. With the ultimate roundup, we'll be looking at ALL the office Suites, including STAR OFFICE 6, Hancorn Office, OpenOffice.org and more – find out how they perform, and which is the best for you.

## Mandrake 8.2 and Red Hat 7.3

Hot on the heels of SuSE8, boxed versions of the other two big distros have reached these shores – check out who's winning the battle of desktop distros.

## Plus: also on test –

Vmware ESX Server, Win4Lin4, Jrun4, Connectiva 8.0 – the RPM-based distro which updates with *apt*, VIAs latest uATX motherboard and more...

## Including LINUXPRO

The third issue of Linux Pro will feature more real-world Linux, including case studies, security notes, interviews and more!

**DON'T MISS YOUR COPY OF LINUX FORMAT!**

Get it delivered to your door every month – subscribe on page 96

The exact contents of future issues are subject to change

# LINUXPRO

From the makers of LINUX Format

July 2002

## GPL goes to **court**

Will the GPL stand up to legal scrutiny?  
The Free Software licence on trial inside...

### **PLUS:**

**Eclipse** – IBM's open source tools platform investigated

**Tripwire** – Your last line of defence against crackers

## Contacts

### 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

Richard Smedley  
richard.smedley@futurenet.co.uk

#### Photography

Getty Images

#### Publisher

Dave Taylor

### ADVERTISING

#### Group advertising manager

Simon Moss

#### Advertising manager

Matt Stanley  
matt.stanley@futurenet.co.uk

#### Deputy advertising manager

Emma Mimmack  
emmamimmack@futurenet.co.uk

LinuxPro is a part of



magazine

Published by



Media with passion

30 Monmouth Street,  
Bath BA1 2BW  
Tel +44(0) 1225 442244

# Welcome

Twenty-four pages of real-world Linux for IT professionals

**Welcome back!** If you missed our very first issue of *Linux Pro*, then you'll have missed out on some thought-provoking material from our IBM case study, an interview with the president of the Free Software Foundation – Europe, and a look at viruses, and what threats they pose to Linux systems.

Judging by the response we've had so far, it seems that you like this 'mini' magazine, which is good news because here's another one! This time we will be looking at the legalities of the GPL and Free Software licences. It seems a bit unbelievable that the terms of the GPL have never been tested in court, but that situation nearly changed recently when a dispute between MySQL and NuSphere saw them attending the US courts. Hopefully our look at what happened and exclusive interview with MySQL's CEO will help answer your questions.

IBM are back with us this issue, as we



delve into their Eclipse tools to discover what they offer to developers, and in our security focus this month, we'll take a brief look at tripwire.

Do keep your comments and suggestions coming in – just email me at the address below!

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

## Contents

**THE GPL IN COURT:**  
Licence legalities on trial **p4**

**ECLIPSE:** IBM's open development tools explored **p11**

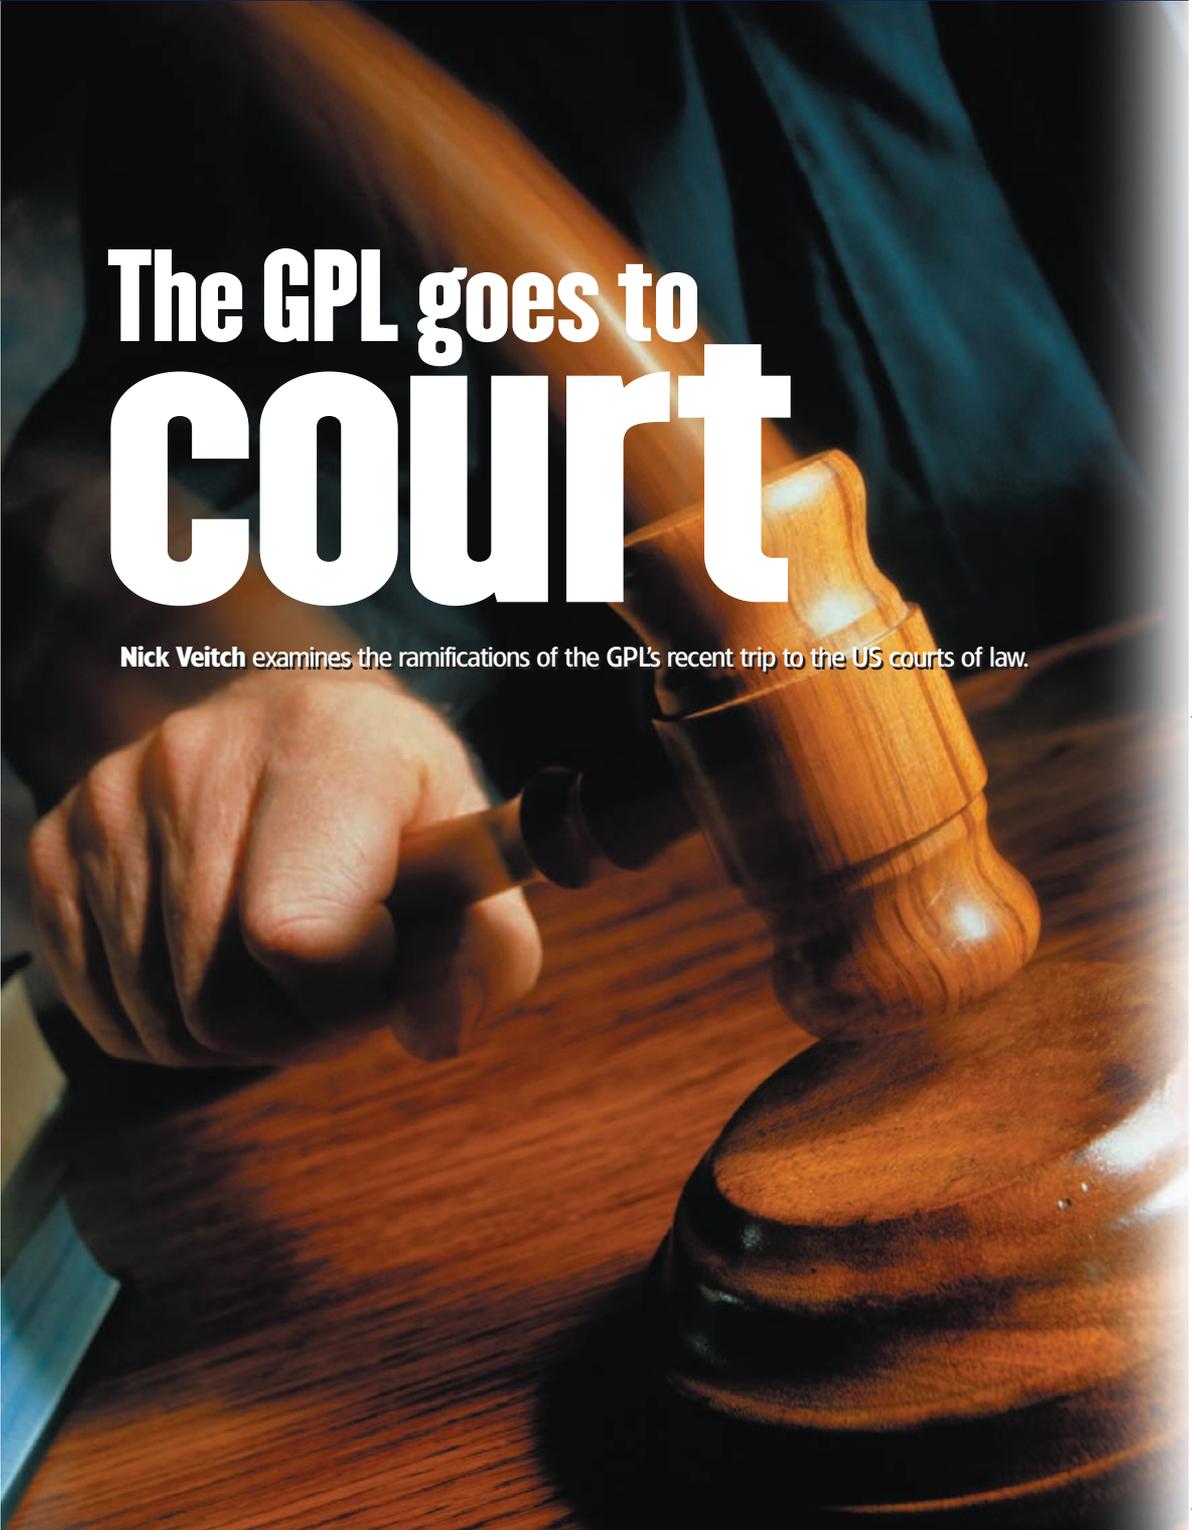
**TRIPWIRE:** Your last line of defence against attack **p20**



**FEATURE | GPL IN COURT**

# The GPL goes to **court**

**Nick Veitch** examines the ramifications of the GPL's recent trip to the US courts of law.



**F**or almost as long as anyone can remember, before the term 'Open Source' was even coined, Free Software has usually meant software covered by the GNU General Public License. As opposed to the sort of licences you might be used to seeing in commercial software, like Microsoft's EULA texts for example, the GNU GPL is there to grant you rights, not to take them away.

The existence of the GPL is pretty much fundamental to what people call 'open source', though other licences can be and are used (see the *Other licences* box). For years, the GPL licence has been used to safeguard the freedoms and rights of software users and developers, but to date, the legality of this user licence has never been tested in a court of law. However, a recent legal case between Nusphere and MySQL AB, the company behind the development of the GPL'd MySQL database software, exposed the possibility that the GPL may be used as a legal tool. We'll come to the specifics shortly, but first, to understand the nature of the case, you need to understand a bit about the GNU GPL licence.

### Freedom as defined by the GNU GPL

Free software is distinct from 'Open Source' in general as it often offers more freedoms than other licences of its type. The following text, from the GNU homepages, defines the freedoms it is trying to protect.

**Free software is a matter of the users' freedom to run, copy, distribute, study, change and improve the software. More precisely, it refers to four kinds of freedom, for the users of the software:**

- The freedom to run the program, for any purpose (freedom 0).
- The freedom to study how the program works, and adapt it to your needs (freedom 1). Access to the source code is a precondition for this.
- The freedom to redistribute copies so you can help your neighbor (freedom 2).
- The freedom to improve the program, and release your improvements to the public, so that the whole community

**benefits. (freedom 3). Access to the source code is a precondition for this.**

© Free Software Foundation, Inc.

Those unfamiliar with the concepts behind Free Software often confuse 'free software' with Public Domain, but the two are not the same and have some very serious differences. Free Software does not mean the software has to be distributed at no cost – 'free' is used in the sense of 'liberty' rather than to mean 'something for nothing'. Public Domain software contains no conditions on use. PD software and its source, can and have been used to build proprietary software – the original authors maintain no ongoing rights. The freedoms expounded above guarantee you freedom to sell the source, redistribute it and modify it, but the remainder of the GPL maintains that such modifications also have to be made publicly available under the same conditions.

This maintains your rights to use the code, but makes sure that the resulting software does not end up being 'closed-source'. In fact the GPL tends to convert software the other way, as proprietary software merged with GPL code becomes 'opened'. It is widely believed that it was specifically the GPL licence that Microsoft's CEO Steve Ballmer was referring to when he said "Linux is a cancer that attaches



MySQL – now to be found at .com and .org domains.





**Marten G. Mickos**  
MySQL CEO and GPL  
advocate.

**“We consider the GPL to be stronger now that we have taken the issue to court”**



itself in an intellectual property sense to everything it touches” in an interview with the Chicago Sun-Times. This may seem like an extreme view, but in some quarters such freedom is apparently considered evil and “un-American”.

Does this prohibit companies from making money out of free software? Not at all. The

examples are too numerous to mention, but to return to the specific case of Nusphere and MySQL, here are two companies that were/are making money from GPL'd software.

### The MySQL story

The collaboration between Nusphere and MySQL goes back to the original open sourcing of *MySQL* in June 2000. Nusphere created a product called *NuSphere MySQL Advantage* based on *MySQL*. This included the main *MySQL* code, provided as source and compiled, including their own *Gemini* code (a transactional storage engine). The problem is, although the binary version distributed contained the *MySQL* software statically linked to the *Gemini* code, the *Gemini* source was not included on the disc.

This contravenes the GPL – effectively, Nusphere had used the *MySQL* to create a ‘closed’ or proprietary product. Often such instances are usually cleared up by the infringer being informed that they have broken the licence terms and that they either have to GPL the whole software, or withdraw it. The situation got more complicated as Nusphere also registered the [www.mysql.org](http://www.mysql.org) domain and published a website which failed to credit MySQL Ab as the creators of the *MySQL* software. They also continued to publish the *MySQL Advantage*, although agreement to use the *MySQL* trademark had effectively expired.

### In court

The initial hearing was held at the end of January. To the disappointment of any of the people who had travelled to the court with interests in the GPL (including FSF's Eben Moglen), Judge Saris refused to hear any arguments based on the on the grounds that the trademark infringement alone was strong enough.

In some ways this is a bit of a let down for the FSF legal team, who no doubt were looking forward to being able to enforce the terms of the GPL for the first time. There is still some doubt in legal circles whether ‘copyleft’ (the GPL's phrase for its copyright terms) is actually an enforceable copyright, or merely a



contract, like most other software licences. However, the lack of any legal enforcement is seen by some as evidence of the great strength of the GPL – most infringements are settled out of court.

But do businesses feel safe in protecting their code with the GPL, and is some uncertainty a price worth paying for the advantages of free software development? We could think of one good person to ask, the CEO of MySQL, Marten G.Mickos:

**LinuxPro:** *MySQL* is widely regarded as the first popular database system to be released under a GPL licence, but it wasn't originally released under that licence. Can you explain why *MySQL* GPL'd 3.23.19 and subsequent versions of the software?

**MM:** Our source had always been open and we were looking at adopting a generic licence that others used as well. It became apparent that GPL would be the best licence for us as it is good at protecting intellectual property and trademark rights while at the same time making it very easy for users to use and adopt the software. GPL also was becoming so popular that nobody had any questions about it. If a software was GPL, people knew how they could use it, and Linux distributors knew they could easily include it in their box.

**LXP:** At that time did you consider that publishing under the GPL might have any legal ramifications?

**MM:** We were definitely aware that ANY licence text may have legal ramifications, but we did not expect to be involved in a legal proceeding about the GPL one year later. (We actually still do not really understand why Progress decided not to initially honour the GPL license.)

**LXP:** Do you still consider the GPL as a robust licensing model?

**MM:** You might say that we consider it stronger now. We have taken the issue to court, and that fact in itself (without a final ruling) has already induced the other party to rectify its non-compliance.

**LXP:** Do you think that any terms of the GPL are unclear – is it possible that people might unwittingly breach the terms of the licence?

**MM:** It is possible to misunderstand the GPL, or not to see all characteristics of it, and it is still a rather new licence text for most people. So in that respect it is natural that people misunderstand it and unknowingly breach the terms of it (but I would say that this is typical of any licensing text). But in nearly all cases, a



The MySQL team bought into 'Free Software' just a year ago, but have already reaped huge benefits.



## FEATURE | GPL IN COURT

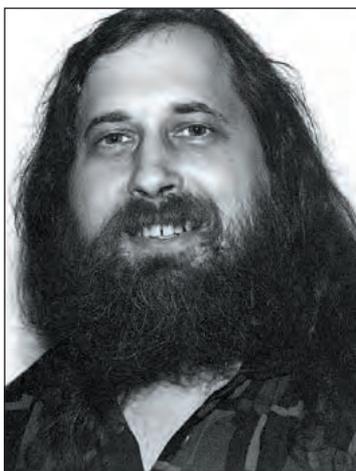


simple and polite communication with the potential breacher has cured the breach, if there were one. So this is not a big problem for us. On the contrary, this is a good starting point for productive communication between the parties.

**LXP:** Do you think there are any terms in the licence that would have made your case against NuSphere more clear cut. There has been some discussion about the termination clause in section 4, and whether violation means permanent loss of licence?

**MM:** In our legal battle the other side has not presented any objection to section 4, so I believe we are in a rather good agreement on what it means. Some people have asked whether section 4 would apply to ANY software under GPL, but our interpretation is that it cancels the GPL only in respect to the product whose GPL was breached. In other words, breaching the GPL of *MySQL* means that you lose the right to use *MySQL* under GPL, but your right to other GPL'd software remains intact.

**LXP:** We understand that Nusphere previously tried to sue MySQL AB and some of its employees personally, prior to the recent trademark dispute. Is this not something that could have been settled out of court?



**Originator of the GNU Public License and father of Free Software, Richard M Stallman.**

**MM:** Yes, and we tried! When it became apparent that NuSphere had acquired the mysql.org domain, we immediately jumped on aeroplanes from Sweden and Finland and flew over to them to discuss and understand what was going on. It was at the end of that three-day meeting that they offered yellow envelopes to us with the summons, instead of a hand to shake. I tried to call the president of NuSphere that same evening to understand what was going on (as he was not present when they sued us), but he never returned my call. Then we presented a public invitation to NuSphere to come to London for settlement talks, but they chose to do a two-hour conference call where they started by saying that they had no authority to agree on a settlement. And so it went on.

**LXP:** Is your experience with Nusphere likely to stop you entering into similar agreements with others?

**MM:** Not at all. Misfortunes happen, but you have to get over it. We have signed numerous successful partnership agreements since. We are now more attentive to what rights we grant and not grant and how those things are presented in the contract text, but that is all.

**LXP:** Is the loss of Nusphere revenue likely to impact on your finances?

**MM:** I'd say loss of revenue always impacts finances, so in that sense the answer is yes. We should have received payments from NuSphere that we never received, and we also have had extra costs for the lawsuit. But at the same time we have developed our other business, and we are fortunate to have a very strong business situation at the moment. Sales in 2002 is more than 50% over budget, both for licences and support/services.

**LXP:** Are the mechanics of an open source company like yourselves, entering into agreement with a "closed" company, necessarily more complicated than for two proprietary companies dealing with each other?

**MM:** Doesn't have to be, but understandably MySQL was a rather thinly organised company





**“We are now more attentive to what rights we grant and do not grant, but that’s all”**

when the initial deal with Progress was signed. Monty and David (our founders) had always done successful business by handshake and by trusting and being trusted. They had so many excellent business relationships with people and companies around the world that they did not think of how wrong things can go. And open source also has a taste of “I am being open with you”, so you might not expect the opposite. Naturally many other open source companies have similar backgrounds. Great guys with great products, but with little financial and legal assistance. But I would say the decisive thing is not being open source, but being a not-so-big growth company. Other growth companies (open or closed source) probably have similar situations.

**LXP:** In the injunction hearing, Judge Saris granted the injunction purely on the grounds of trademark infringement, without involving the licence issues at all.

**MM:** Yes. She specifically said that she is not convinced of the “irreparable harm” although she felt that MySQL AB had the stronger case. Technically, if you can cure your breach later on by paying some amount of money, then the harm is not considered irreparable. So that was the reason. (And, as you may know, in US legislation, “irreparable harm” is a key condition for granting a preliminary injunction.)

**LXP:** Aside from the legal aspects of the GPL, do you think that open sourcing your software has been good for business.

**MM:** Fantastic! It is a new way to produce and to distribute software. The database business has not seen a new major player in the last 15 years, but now there is one. Of course the product has to be superb in order to succeed, but when you enter a mature market (like we did when we entered the database market), you need a new business model as well. (Look at DELL as an example of a similar situation, but in the PC market).

**LXP:** Do you receive many external code contributions for the core database server, and if so, how to you handle the copyright of these?

**MM:** Not many, but some. In order to include them in the product, we ask the contributor to share the copyright with us. This means that we get full copyright and ownership of the code, but the original contributor retains his own right to do whatever he wants with the same piece of code.

**LXP:** Many people in the business community tend to view “open source” and the GPL as a strange idea, and can’t see why they would want to effectively “give away” their intellectual property. Do you have any words of wisdom for these people?



## Other Free Software licences, and some not so free...

To confuse things further, there are many sorts of software licences which describe themselves as free or 'open source'. Some of these have evolved because they are built on other code which had different restrictions.

Others make special conditions which are incompatible with the GNU GPL. Most free software is covered under the GPL and compatible licences, but some are not. Usually, incompatible licences fail to deliver one or more of the freedoms outlined, which doesn't necessarily make them incompatible with the ideas of 'Open Source'. This is not an exhaustive list. You can find out more from [www.gnu.org/licenses](http://www.gnu.org/licenses).

### GPL compatible licences

**GNU LGPL** – The Gnu Lesser General Public License permits linking to non-free modules, and mainly grew out of the need for libraries to interact with

other software.

**The X11 licence** – a simple licence used by XFree86 and other projects.

**Modified BSD licence** – The original BSD licence, with the clause about advertising removed

**W3C Software Licence** – a licence defined by the W3C(WorldWideWeb Consortium) to cover software used/developed for web standards.

**Intel Open Source Licence** – which covers a lot of Intel's free tools

### GPL incompatible licences

**IBM Public Licence** – Covers a lot of IBM software but is incompatible with GPL because of patent clauses.

**Mozilla Public Licence** – original draft not completely compatible with GPL, and imposes restrictions that make it hard to work with GPL software

**Aladdin Free Public licence** – imposes restrictions – isn't really free at all.



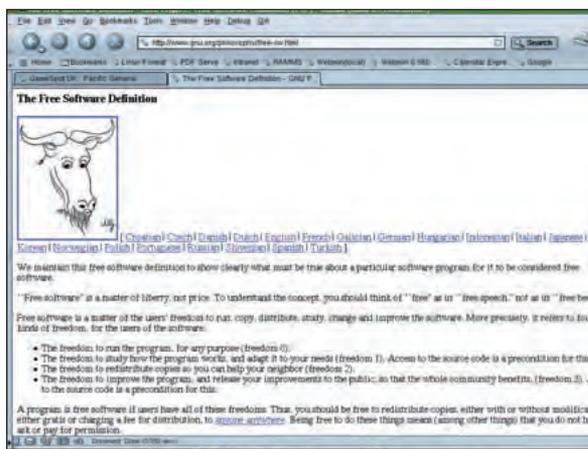
**MM:** I would start by reminding them that all source used to be open – in the early days of computing. Then, later, people started to hide their source. The question should be “Why

would anyone want to hide his source code?” Copyright law gives full legal protection of your intellectual property.

Of course it is easy to copy open source code, but that is illegal. It is also easy to copy closed source code, and that is illegal as well. So the legal protection is not so much better for closed source (but, admittedly, there is a theoretical risk of someone using somebody else's open source code in his own closed source product, and that might be difficult to detect).

In our case our customers are so dependent on our skill to develop the software and correct any bugs in it that they come to us anyhow. And the fact that the source is open is the best guarantee that there are no hidden loopholes in the software.

When I meet people who are in doubt, I also show them all the big corporations who endorse or use or produce open source software (investment banks, IBM, HP, etc. ). ■



You can find out much more about Free Software at [www.gnu.org](http://www.gnu.org)



# fits all SIZES

Richard Drummond  
examines IBM's Eclipse –  
an open tools platform,  
extensible for any task.





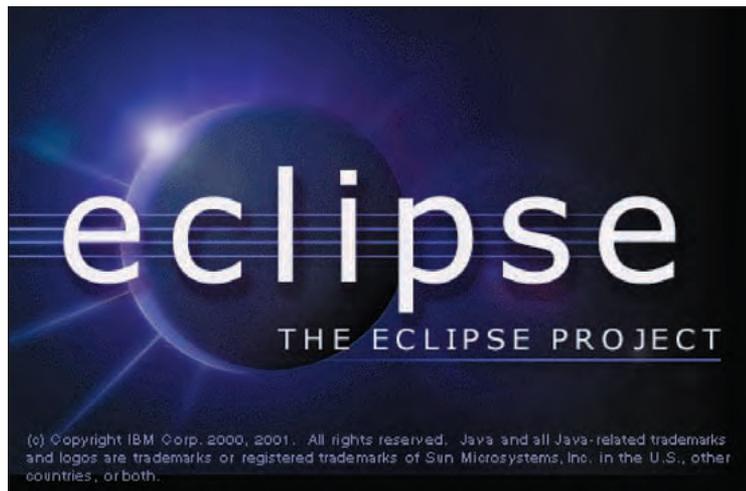
**A**pache, some would say, is Linux's killer application. IBM recognised this some time ago and turned around its moribund and confused application server business, basing its range of *WebSphere* server products on the solid, open source foundations of the *Apache* web server. However, while Linux and open source currently have the deployment of e-commerce applications pretty much sewn up – with technologies such as *Apache*, *MySQL*, *Tomcat* and *JBoss*, for example – the open source arena has lacked a compelling and integrated solution for the development of web-based and enterprise-level applications. For those used to Unix, console-based tools such as *Emacs*, *make* and the *GNU Compiler Collection* are powerful weapons; but for those used to the point-and-click, rapid-

development nature of *VisualBasic* and similar products, the traditional Unix tools can cause something of a culture shock. Enter Big Blue to the rescue once more. Its next-generation *WebSphere Studio Application Developer*, a leading-edge visual J2EE development environment, is also now based on an open source foundation: Eclipse.

### What is Eclipse?

The Eclipse project was initiated by IBM in November 2001 with the donation of an estimated \$40 million of source code. Its aims are simple: to create an open, robust, commercial-quality development platform into which commercial vendors can plug in their own tools to gain interoperability. Eclipse provides a framework which can unite source-code editors, source-code management tools, compilers, debuggers and CASE tools within a common integrated environment. The Eclipse code is licensed under the OSI-approved Common Public License, and can be redistributed worldwide without restriction. This license allows vendors to build their own commercial, branded products on top of the common Eclipse platform and to supply commercial tools as plug-ins to extend the capabilities of Eclipse.

**“An open, robust, commercial-quality development platform into which vendors can plug in tools.”**



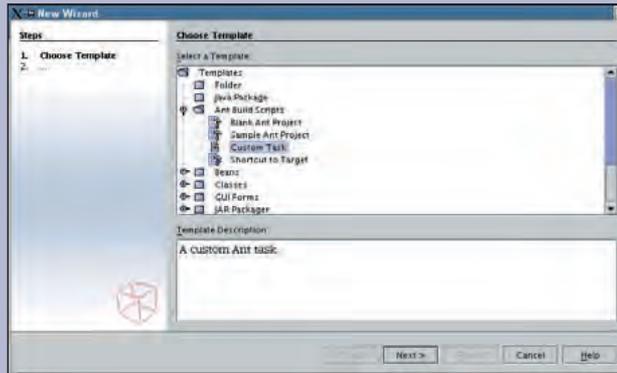
Eclipse: a new dawn for multi-language, open platform development?



## Eclipsing Sun? Open architecture rivals

One obvious omission from the Eclipse Consortium, with the exception of Microsoft, is Sun Microsystems. This snub is made doubly worse when you consider that Sun already fund an open source project with an architecture similar to Eclipse. A year before Eclipse began, it donated the code to its Forte integrated development environment to the open source community and set up the NetBeans project (see [www.netbeans.org](http://www.netbeans.org)). Like Eclipse, NetBeans is a modular, Java-based tools framework.

NetBeans does have a head start on Eclipse. It is a more mature, better understood and better tested project. It has a larger open source developer community, more advanced tools in the open source domain, and arguably more third-party development; in particular, NetBeans offers better integration with key open source technologies, with plug-ins available for development with *Ant*, *Tomcat* and *JBoss*, for example. Another point in NetBeans' favour is that it offers a visual form builder for generating GUIs with *AWT* and *Swing*, not something that has yet been addressed by Eclipse. However, IBM have been more successful, or at least more vocal, at



### Is Eclipse a blatant imitation of Sun's NetBeans?

signing-up commercial partners to the Eclipse platform than Sun have with NetBeans. Eclipse too seems to focussing more on multi-language development: the Eclipse C/C++ development plug-ins, for example, are at a more advanced stage than similar tools for NetBeans.

Supporters of NetBeans, especially Sun, have been quick to cry foul, and accuse IBM of deliberately forking the open source community with the release of Eclipse. Worse, they point out that Eclipse introduces a third, non-standard GUI API for Java rather than working within the existing *AWT* or *Swing* frameworks. You can see

their point. But is it necessarily wasteful to have two open-source projects with similar goals? The reverse has typically proven to be the case. A healthy air of competition and rivalry between NetBeans and Eclipse will help motivate developers and drive forward the progress of both projects, aided by a cross-fertilisation of ideas between the two. Ultimately, two open-source IDEs means more choice for the user, which is always a good thing. Now, if only Sun could only get over its sulks and collaborate with Eclipse to help engineer some degree of interoperability between the two platforms...

Although Eclipse was born out of IBM, the Eclipse project is an independent entity, guided by a consortium of tools vendors and the open source community – united as [eclipse.org](http://eclipse.org). The member list of the Eclipse Consortium reads like a *Who's Who* of development tools suppliers. Alongside IBM, the founding members were Borland, Merant, QNX Software Systems, Rational Software, Red Hat, SuSE, TogetherSoft and WebGain; more recent additions to the board include Serena, Sybase and Fujitsu. Currently, IBM and

Rational Software ship tools based on the stable version of Eclipse platform. Expect more vendors to follow suit when the next generation of Eclipse – version 2.0 – ships later this year.

The Eclipse platform itself is implemented in Java, but is not restricted to providing only Java tools. Plug-ins can be created to support development in any language, including C/C++, Perl, Python and Javascript. Nor is Eclipse limited to only server-side development: anything from the server to the





desktop to the embedded space and in between can be supported with appropriate tools supplied as plug-ins.

Besides the Eclipse framework itself, eclipse.org maintains several sub-projects which make up the functionality of the open source Eclipse download: this currently includes the Platform, JDT, and PDE projects. The Platform project provides the common services – such as the developer desktop itself (called the Workbench), resource management, and language-neutral access to source code repositories – that are available to all Eclipse environments; the JDT or Java Development Tools project adds tools for Java development within Eclipse; and PDE is the plug-in development environment, which supplies tools to speed the creation of Eclipse

plug-ins. A new project, CDT or the C/C++ Development Tools project, was launched in January of this year under sponsorship from Red Hat and IBM. It is not included with Eclipse 1.0, but should be stable in time for the release of Eclipse 2.0. With the JDT and CDT modules installed, it is possible to do Java and C/C++ development side-by-side in Eclipse within a single, integrated environment.

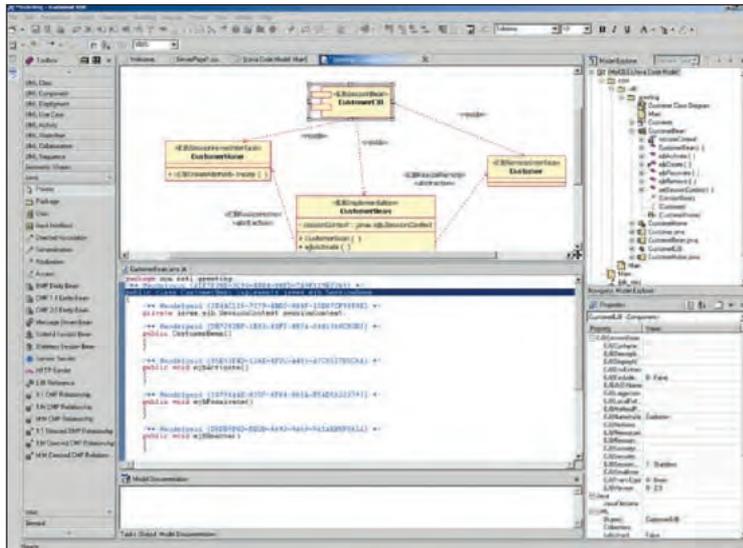
### The open source advantage

IBM's donation of the Eclipse source code isn't motivated by altruism. Far from it. It's part of a clear middleware strategy to supply tools and platforms to enable e-commerce applications. While you can download and use the open-source Eclipse platform for free, IBM hope that the proprietary tools, packaging and support that they add will persuade you to buy their products. So, why have they taken the open source road with Eclipse?

Besides all the usual benefits of open source development, such as stability and rapid evolution, the key advantage that Eclipse brings is interoperability. Basing Eclipse on open standards and publishing its source-code and APIs openly allows anybody to take part and to provide tools that will integrate with

**“Developers have a wide choice of tools that can be assembled into a seamless integrated environment”**

Eclipse can support tools to cover the entire software cycle – behold Rational's UML tools.



## The third way The Standard Widget Toolkit

The existing Java APIs for GUI development, *AWT* and *Swing*, have their fair share of detractors. *AWT*, they say, is too simple, while *Swing* is too complex and too slow. Reasons such as these have led to the creation of third Java GUI toolkit for use in the Eclipse platform, the *Standard Widget Toolkit* or *SWT*. In many ways, *SWT* is a half-way house between *AWT* and *Swing*. Like *AWT*, it provides a common, platform-independent Java interface for native widget sets such as *Motif* and the Windows GUI; unlike *AWT*, it aims higher than the the lowest common denominator. Where a widget is unavailable on the underlying GUI platform, it is emulated by *SWT*. Moreover, *SWT* is efficient, with a performance close to that of the native platform, and offers a native look and feel. (A complaint against *Swing* is often that developers, especially Windows developers, don't want their applications to look like a Java application.

Many claim, unfairly in my opinion, that the default *Swing* theme is ugly and neglect to notice the fact that *Swing* is skinnable. Motif: now that's ugly).

The problem with *SWT* – apart from the fact it introduces a third, incompatible GUI framework – is that it breaks Java's write once, run anywhere philosophy. Unlike the pure Java *Swing* toolkit, *SWT* relies on a native layer, and thus requires porting to any new platform. Currently, Windows and *Motif* on Linux are supported, and work is under way to create ports for the *GTK+* toolkit on Linux, *Motif* on Solaris, AIX, and HP/UX, *Carbon* on MacOS and *Photon* on QNX.

*SWT* is also licensed under the Common Public License and can be used in your own Java projects; it isn't restricted to use just within Eclipse. The Eclipse.org web site host tutorials and documentation to help developers learn to program with *SWT*.

Eclipse. The immediate advantage for the consumer (the developer) is that he or she has a wide choice of tools that can be assembled into a seamless, integrated environment suitable for the job at hand. If you start a new project requiring different tools, you don't need a whole new development platform; you just have to obtain any additional tools you need as Eclipse plug-ins. Because of the common framework that Eclipse provides, tools running under Eclipse will share a common look and feel, and so less retraining will be required than if you had shift to a totally new development platform.

Perhaps the most compelling reason for to adopt the Eclipse platform, however, is choice: developers are no longer locked in to choosing tools from a single vendor. A developer could pick an editor from company A, a compiler from company B, and debugger from company C, all based on the merits of the individual products, confident that – as long as these tools are available as Eclipse plug-ins – they will work together.

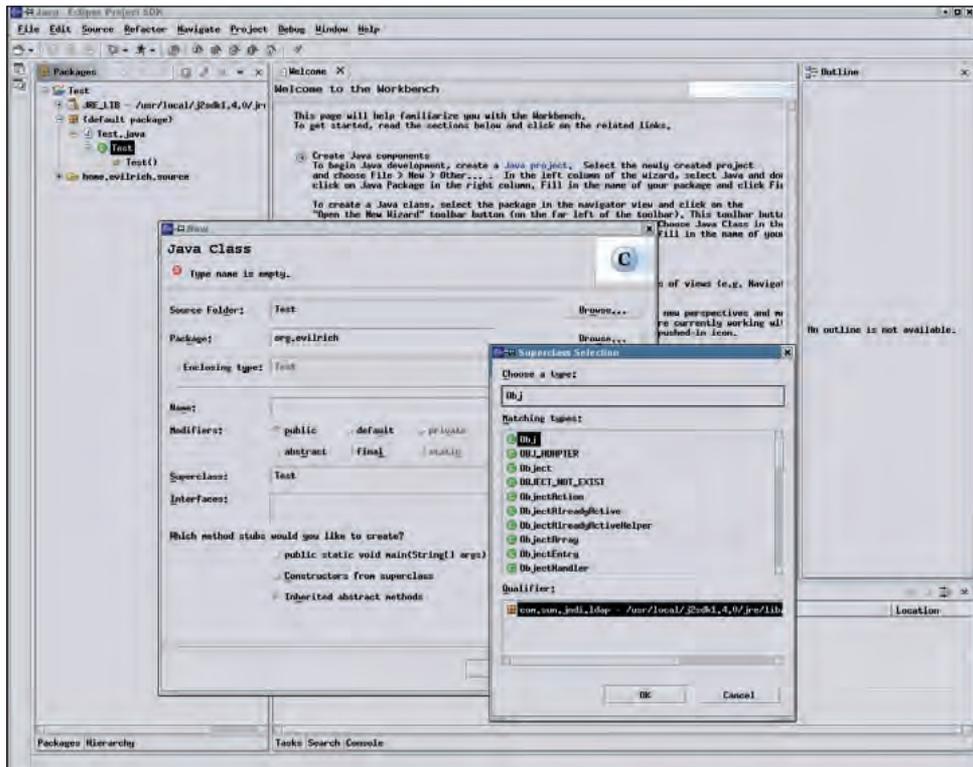
IBM have realised that the can't win the war against Microsoft by creating yet another proprietary development platform. However, by opening up the Eclipse platform, by leveraging tools from other vendors, and by providing the developer with choice, they can offer a product which is more attractive than anything Microsoft can offer. Or that's the plan.

### Architecting openness

Eclipse's strength is derived from its open, modular architecture – so it's worth taking a closer look at this. Eclipse is built on a system of plug-ins written in the Java language and typically supplied as JAR files. Everything bar the core runtime – including the Platform components – is provided as a plug-in. When the system is launched, it discovers what plug-ins are installed, and for each plug-in finds what services it provides and how it is interconnected with the other available plug-ins. A registry of plug-ins is built in memory, and plug-ins are only started when their services are required, thus speeding up the



## FEATURE | ECLIPSE



**The Eclipse workbench can be extended to provide language-specific environments. This is the Java project.**

◀ startup time of Eclipse as a whole. When you start Eclipse normally, the Platform components are loaded. These provide the Workbench, the graphical desktop upon which the other plug-in tools and editors can embed language-specific views and editors in a manner that will be instantly familiar to anybody who has used an IDE before. The user interface of the Workbench is built upon libraries supplied by Platform itself. The basic widget set is provided by the *SWT*, the *Standard Widget Toolkit*, a new GUI toolkit for Java that implements an efficient, OS-independent layer for native GUI toolkits (see box, *SWT: The third way*). Layered over *SWT* is *JFace*, a library which furnishes higher-level interface elements such as wizards and views. On top of that pile is the Workbench itself. Of course, since all the components of the Eclipse GUI are made available to the system

as plug-ins, you don't have to make use of them at all if you wish – for instance, if you were to build a console-based application on top of the Eclipse framework.

The Platform sub-project, if you remember, provides a set of core services which tools vendors can add to and extend to build their own tools and environments tailored for specific languages or tasks. It offers no language-specific tools by itself. Besides the Workbench framework, Platform includes the infrastructure for debugging, for the online help system, for scripting and for concurrent version management. It also supplies a generic search component and a component for WebDAV integration.

The JDT sub-project builds on the Platform project to provide plug-in tools that implement a powerful Java programming environment. JDT extends the Workbench with views,



editors and wizards targeted for Java development. This includes a Java source code editor with syntax-highlighting, code-completion and code formatting; project and source code navigation tools; and code refactoring tools. The JDT also features an incremental Java compiler, and provides integrated facilities for launching and debugging Java code. Of course, the JDT also allows Eclipse to be a development environment for itself, and this is assisted by the PDE or Plug-in Development Environment which builds on the JDT to add views and editors tailored for building Eclipse plug-ins.

Tools vendors can extend the JDT plug-ins themselves to add more features to the Java development environment – and this is what IBM have done with their *WebSphere Studio Application Developer* which enables J2EE and Web Services development, and extends the JDT project's capabilities with editors, wizards and tools to enable the rapid development of JavaBeans, Java Servlets, JSPs (Java Server Pages) and EJBs (Enterprise Java Beans).

As we have said, however, Eclipse is not limited to Java development, nor to simply the implementation and testing phases of the software lifecycle. The CDT (C/C++ Development Tools) project, for example, is working towards providing a fully functional C and C++ environment for the Eclipse platform. The current focus is C/C++ development on Linux for deployment on Linux, but support for other platforms including Windows and embedded applications are planned. CDT extends the core Platform components with a source editor for C/C++ code, navigation tools, and an integrated debugger. Another sub-project hosted by Eclipse.org, but not much beyond the planning stage, is the Graphical Editor Framework (GEF) Project. This aims to create a generic or application-neutral graphical editor which may be used by tools vendors as a basis for UML diagram editors, GUI builders or even WYSIWYG text or HTML editors.

### Getting commercial

Open source doesn't necessarily mean free, and the whole purpose of Eclipse was to

provide an open basis for IBM's commercial development tools, while leveraging the technology and talents of third-party tools vendors. Using Eclipse as a foundation, IBM has been able to revamp and consolidate its range of development tools, and, now, as with its server platforms, it will market them all under the *WebSphere* brand. So, gone is the old *VisualAge* for Java product line and out is the old *WebSphere Studio Professional* and *Advanced* editions. The new Eclipse-powered *WebSphere Studio* product range are all based on the Eclipse Platform – which IBM are calling *WebSphere Studio Workbench*. Initial products in this stable include *WS Site Developer* and *WS Application Developer*. *Site Developer* is aimed at teams of developers doing e-commerce development with Web Services, servlets and JSPs, while *Application Developer* is a full-blown J2EE environment. Both products feature plug-ins from third-party vendors, including Rational and Merant.

Rational Software themselves, famous for their leading CASE tools, also ship products



**“Using Eclipse as a foundation IBM has been able to revamp its range of development tools.”**

### Eclipse at SourceForge Alphabet soup in the making

A number of third-party Eclipse plug-ins are beginning to appear as projects on SourceForge, the site which provides hosting services for open source projects, although many of these have not progressed much beyond the planning stage. One such example is the *SolarEclipse* project, which aims to provide XML and web application development tools for the Eclipse platform, including an XML editor, HTML and JSP editors, servlet wizards and plug-ins for integration with the *Tomcat* application server (see <http://solareclipse.sf.net>). Another project is *LunarEclipse* (<http://lunareclipse.sf.net>), who are developing an RMI plug-in for Eclipse which can generate RMI stubs and skeletons. Why not pop along to <http://sf.net/> and search for projects containing the keyword 'Eclipse'?





based on Eclipse. *Rational XDE Professional: Java Platform Edition* builds on Eclipse Platform to create a Java environment with tools to accelerate both the design and development of J2EE solutions. *Rational XDE* includes a powerful, customisable pattern engine with nearly 60 patterns that can be used as a basis for your J2EE solutions. Also included is *InstantUML*, an integrated UML design tool. *Rational XDE* may also integrate with other Rational tools, such as *RequisitePro* and *AnalystStudio*, for requirements management and analysis. Rational also supply plug-ins for Eclipse to integrate with its *ClearCase* software configuration and change management tools.

An interesting Eclipse-based product currently shipping is LegacyJ's *PERCobol* suite. The *PERCobol* compiler allows COBOL software to be compiled using Java Technology and to build and deploy COBOL apps as Java applications, servlets, Java Beans or EJBs. It

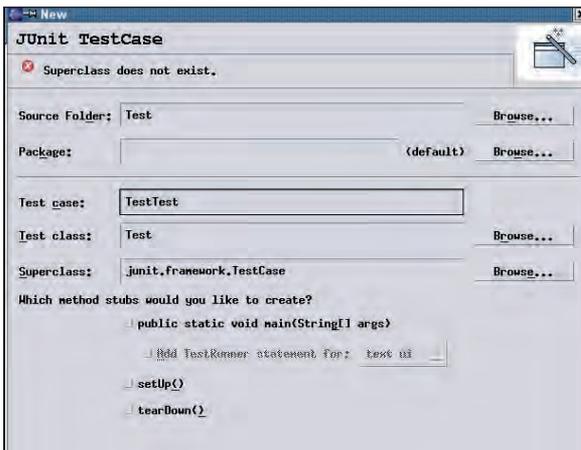
includes support for XML and SOAP, thus enabling you to migrate existing COBOL software into Web Services.

## Getting Eclipse

The availability of commercial tools will be key to the success of Eclipse, but one benefit of its open source architecture is that you can download and use the basic platform – complete with tools for Java and C/C++ development – for free. This may not have much appeal for serious J2EE developers, but it is certainly a significant plus for smaller software shops, students and the enthusiast. The open source Eclipse currently doesn't offer much support for J2EE development, XML or Web Services – but it is more than adequate as a general IDE. Expect more advanced tools to appear as open-source plug-ins as the Eclipse community grows.

You can download source and binaries for Eclipse at [www.eclipse.org](http://www.eclipse.org). Currently, the initial version 1.0 release and various development snapshots of version 2.0 are available for Windows and Linux (version 2.0 builds are also available for the platforms targeted in version 2.0, such as *GTK+/Linux*, *Motif/Solaris* and *Carbon/MacOS X*). The developers issue several series of Eclipse 2.0 builds. Nightly snapshots as well as integration (testing) builds are available, but, unless you are an Eclipse hacker, we recommend you stick to the stable, milestone builds. Make sure you download the full Eclipse SDK (which includes the Eclipse Platform, JDT and PDE) rather than just the Platform alone – the Platform isn't much use by itself. The CDT (C/C++ tools) project is available as separate download. Get the one that is targeted for your build of Eclipse.

**Eclipse 2.0 offers new features, such as support for the JUnit framework.**



**“It is quite possible that IBM has another open source killer application on their hands”**

## Looking forward

It is impressive to see how far Eclipse has come in so short a time. It's almost hard to believe that the open source project was launched only eight months ago. Eclipse version 2.0 is scheduled to be released at the end of June and will add many new features and ports to new platforms. As well as Linux and Windows, the primary targets for version 2.0 will be Solaris, HP-UX and IBM's AIX. A



## The eclipse.org community A development focus

Any source projects needs a website – to act both as a hub for developers and to provides a point of focus and resources for developers and users alike. Eclipse is not different.

At [www.eclipse.org](http://www.eclipse.org) you can find all the usual material: news, downloads, mailing lists, documentation and a bug tracking system. The news section is handy for those wishing to keep abreast of Eclipse developments, since it covers changes within the open source projects and keeps tabs on any commercial projects based on the Eclipse framework. The documentation section is also good, hosting user manuals, specifications, developer documentation and tutorials. If you want to find out how to make your own Eclipse plug-in, for example, or how to developer with the SWT in your own projects, then you'll find the information to help here.

One omission from [eclipse.org](http://eclipse.org), however, is any facilities for hosting third-party open source



**Eclipse.org: every community needs a home.**

projects to develop Eclipse plug-ins. In this regard, it lags behind the NetBeans community site. If you are

developing your own plug-in, you'll have to pursue existing channels such as SourceForge for hosting.

whole raft of improvements are promised, including various usability tweaks to the Workbench; better integration with *Ant* for project building; and improved support for accessing CVS repositories. The JDT project will see enhancements too, such as a JDK1.4-compliant compiler, more refactoring tools, Javadoc generation and support for JUnit unit testing framework.

Although version 2.0 promises some real improvement, the real test for Eclipse is how well the C/C++ tools will be received. Linux developers have been calling out for a powerful IDE for C/C++ development on Linux, and the Eclipse CDT project will be crucial in the acceptance of Eclipse by Linux developers. If it proves popular, then expect

Eclipse plug-ins for developing in Perl and Python to follow, truly transforming Eclipse from a Java IDE to a multi-language platform.

It is hard to see – with the backing of currently around 30 of the leading tools vendors – how Eclipse cannot be an immediate commercial success. In the unlikely event that it is not, because of its open source foundations, your investment in Eclipse technology will not be wasted. What's more, the open source Eclipse project itself will become more valuable as new features are added, as the developer community grows, and as more third-party plug-ins are developed. It is quite possible that IBM has another open source killer application on their hands. ■



**TUTORIAL** | TRIPWIRE

**INSTALLING AND RUNNING**

# Tripwire



In this month's security tutorial Charlie Stross explains how to install Tripwire and tell when the bad guys have done something naughty to your server.

**T**ripwire is the last line of defence against hacking or unauthorised attacks on your computer. If the bad guys can get through your firewall and get shell access to your machine, then generally they do one of two things: crash or trash it, or install their own files and use it as a springboard for attacks on other computers. A typical modern Linux workstation can have a quarter of a million files on it – KDE 3, alone, has over 16,000 components – and there's no way you can spot when a corrupt binary has appeared, or a configuration file has been subtly edited, without outside assistance. Enter *tripwire*.

*Tripwire* is a policy-driven filesystem scanner. That is, it has a set of configuration files and a database that tells it what your filesystem should look like. (You can exclude those chunks of the filesystem that aren't important to system integrity, such as individual user accounts: serious hacking attempts focus on subverting the operating system tools, and user data is rarely a target.) When you first run *tripwire*, you create a baseline database that contains a snapshot of your important files (and an MD5 (or other) checksum for each of them). You then place the baseline database on a read-only medium (such as a write-protected floppy disk or a CD), so that if a hacker gains access to the machine they can't spoof the database. Subsequently, you run the *tripwire* tool to scan the filesystem, and it will report only those files that have changed. If you are confident that the machine is secure and the only changes are official ones, you can merge approved changes into a new baseline database copy: if you've been attacked, the

output from *tripwire* will tell you what files the attackers have changed.

*Tripwire 1.x* was originally a free software project, released under GPL, and written to run on any Unix. The developers subsequently created a company, Tripwire Security ([www.tripwire.com](http://www.tripwire.com)), to sell commercial versions of *tripwire*, numbered 2.x and 3.0 (current). They only really support *tripwire* on Intel-based Linux platforms at present, and they release a free (GPLd) version of the commercial product. (They also sell a very nice graphical management tool called *Tripwire Manager* that enables you to run *tripwire* remotely on a whole network of servers, and detect network-wide attacks.) While you might want to pay for commercial support if you are trying to secure a company network, the GPL version of *tripwire* is still available for free from <http://sourceforge.net/projects/tripwire/>.

## Installing tripwire

*Tripwire* comes packaged with most modern Linux distributions (including SuSE, Red Hat, Mandrake, and Debian). Note, however, that the installers won't set *tripwire* up to work effectively with your system – they'll usually provide a generic set of configuration files and nothing else. Also note that *tripwire 2.3* is only suitable for Intel Linux systems, while earlier versions ran on BSD and non-Intel Linuxes, and other Unixes. For this reason, some quite early versions of *tripwire* are kicking about; for example SuSE 7.3 ships with *tripwire1.2* rather than *tripwire 2.3*.

You can use the older, free, *tripwire*, or download the *tripwire 2.3* sources from <http://sourceforge.net/projects/tripwire/> and compile them yourself. To do this, extract the tar archive, *cd* into the *src* subdirectory, and type **make release**:

```
cd /tmp
tar xvf tripwire-2.3.1-2.tar.gz
cd tripwire-2.3.1-2
cd src
make release
```

The makefile provided with the sources doesn't install *tripwire 2.3*; to do that, you need to edit the file *install.cfg* (in the subdirectory *tripwire-2.3.1-2/install/*), specifying





the directories you want the *tripwire* programs to go into, the various *tripwire* installation options (such as the SMTP host to email reports through), and so on. Then run the script *install.sh*:

#### sh install.sh

This script reads the settings in *install.cfg*, edits the template configuration files that come with the distribution, and installs everything – it also prompts for a passphrase and uses this to digitally sign the configuration file, policy file, and database files (to make unauthorised modification harder).

*Tripwire 1.2* isn't as sophisticated, and if you've installed it from the RPMs you don't need to jump through the above hoops. You can still obtain *tripwire 1.2* sources from [www.cerias.purdue.edu/coast/coast-tools.html](http://www.cerias.purdue.edu/coast/coast-tools.html), and the source distribution contains instructions for building *tripwire 1.2* on a wide range of platforms. (Warning: familiarity with C and *make* would be useful!)

### Running tripwire for the first time

To set up *tripwire 1.2*, you need to create a

**“With a quarter million files there is no way you can spot a corrupt file or binary without assistance.”**

*tw.config* file; this contains a list of directories to be scanned, and the type of changes *tripwire* can safely ignore without reporting to the owner. The resulting information is collected in the *tw.db* database.

Having created a *tw.config* file, you need to run **tripwire -initialize** to build the first *tw.db* file. This will be created in the directories specified to hold databases (at build time), and will be called *tw.db\_hostname* (where *hostname* is the name of your computer). A minimal *tripwire 1.2* *tw.config* file will look like this:

```
/etc +pinugmc1 # be paranoid about
                contents of /etc
```

```
/usr/sbin +pinugmc1 # be paranoid
                  about contents of /usr/sbin,
/sbin +pinugmc1 # /sbin, and so on
/bin +pinugmc1
/usr/bin +pinugmc1
/var/log +pug1 # watch logfiles for
               permissions, owner/group
!/home # don't monitor /home
!/tmp # don't monitor /tmp
```

The structure of this file is described in the *tw.config* manual page.

*Tripwire 2.0* changed a lot of things – *tw.config* is now encrypted, and you can't edit it by hand this way. (An attacker could, in principle, nubble *tripwire 1.x* by installing their own *tw.config* file before the next time *tripwire* ran.) *Tripwire 2.0* stored files in */etc/tripwire* by default. The file *tw.cfg* stores the location of *tripwire*'s data files, while the *tripwire* policy file, *tw.pol*, specifies what system resources *tripwire* should monitor, how the data should be collected, and who should be notified of policy violations. The files *site.key* and *hostname-local.key* are protected using a passphrase, and must be unlocked in order to decrypt the *tripwire* configuration files before virtually any operation – the command-line tool *twadmin* is used to manage these keys. Finally, the *tripwire* databases are stored in a subdirectory called */var/lib/hostname.twd/*, and report files are created under */var/lib/tripwire/report*.

To set up *tripwire 2.x*, you use *twadmin*; you'll need to read the manual page for *twadmin* to get all its options, but for now note that **twadmin --print-polfile** prints the policy file in readable form; you can save this, edit it, and replace the policy file with

```
twadmin --create-polfile policyfile.txt
```

(and *twadmin* will prompt you for the site password before encrypting the new policy file). The *tripwire 2* policy file format is much more complex than that of *tripwire 1.x*, and is defined in the *twpolicy* manual page.

Having specified a policy (or accepted the default one created by *install.sh*), *tripwire 2.x* users should create a default database like this:

```
/usr/sbin/tripwire -m i
Please enter your local passphrase:
Generating the database...
```



```
*** Processing Unix File System ***
### Warning: File system error.
### Filename: /usr/sbin/fixrmtab
### No such file or directory
### Continuing...
Wrote database file:
/var/lib/tripwire/icebook.antipope
The database was successfully generated.
```

Once you have a *tripwire 1.2* database, you really want to move it onto a non-writable medium – a CD-R is a good choice (remember not to make it a multi-session-capable disc!), or a floppy disk with a write-protect tab. The database needs to be mounted somewhere accessible (such as `/mnt/tripwire`) so that when you run *tripwire* it can consult the baseline that records file checksums.

If you're using *tripwire 2.3*, you have less of a problem; because *tripwire 2.3* uses public key cryptography to prevent its databases and configuration files from being written to without a passphrase, it is not possible for an attacker to tamper with it without the passphrase. So as long as you remembered to pick a strong password, you don't need to worry about moving the database to a non-writable medium.

## Tripwire in everyday use

To check your filesystem against the baseline database, and report deviations (such as changes to system configuration files or programs), using *tripwire 2.3*:

```
tripwire --check --email-report
```

(*Tripwire* will write a very verbose report and email it to the designated administrator.)

To check using *tripwire 1.2*:

```
tripwire
```

(*Tripwire 1.2* will make a somewhat less comprehensive report – and doesn't have an option to email it.)

You should add these commands to root's crontab file. To edit the crontab easily, change to the root user and then use the `-e` switch to edit the crontab file:

```
su
```

```
crontab -e
```

*crontab* is a control file that runs specified programs regularly at a set time of day. The

following text, inserted into a crontab file using a text editor, will run *tripwire* once a day at 1am; just delete the line you don't need (depending on your installed version):

```
# field allowed values
# -----
# minute 0-59
# hour 0-23
# day of month 0-31
# month 0-12 (or names, see
# below)
# day of week 0-7 (0 or 7 is Sun,
# or use names)
#
# the following line runs tripwire 1.2:
01 *** tripwire
# and the next line runs tripwire 2.3:
01 *** tripwire --check --email-report
```

## Administering tripwire

You almost certainly want to run *tripwire* daily or weekly from root's crontab file, or log in manually and run it (as above).

As you reconfigure or upgrade your system, *tripwire* will keep warning you about the changes you've made. To get it to stop, you need to merge the changes into the baseline database. Using *tripwire 1.2*, remount the reference database somewhere writable, run *tripwire* with the **-interactive** flag and answer **y** or **n** when it prompts you for each changed file. Under *tripwire 2.3*, use the command **tripwire --update**, which will generate a report, then throw you into a text editor where you can enter X's in checkboxes against each changed item. *Tripwire* will then digest the report and update its database accordingly (as long as you give it the correct password).

In either case, you really ought to read the *tripwire* manpage (for version 1.2) or the *twinro* manpage (for 2.3) and the other pages it refers to. You may also want to buy the O'Reilly book *Incident Response* (by Kenneth R. van Wyk & Richard Forno, pub. O'Reilly, August 2001, ISBN 0-596-00130-4) which discusses tools for responding to hacking incidents (including *tripwire*) at length. *Tripwire* is a big program, and this tutorial barely scratches the surface! ■

