

# A superlative choice

This month we struggled a bit to find superlatives to describe our main feature, not because it wasn't exciting, but because it simply is the biggest, most massive, comprehensive and detailed roundup of Linux Office software ever. Thanks to the recent *StarOffice 6.0* launch (and *OpenOffice.org 1.0*) and the availability of *HanCom Office 2.0*, it also contains the most up to date versions of all the software possible. Considering that most people who use computers spend the majority of their time using common office applications, the capability of Linux solutions to address this area is of critical importance. But whether you are concerned about the viability of Linux for the desktop, or just wondering if you can write a letter to your granny reliably under Linux, you'll find the answer here.

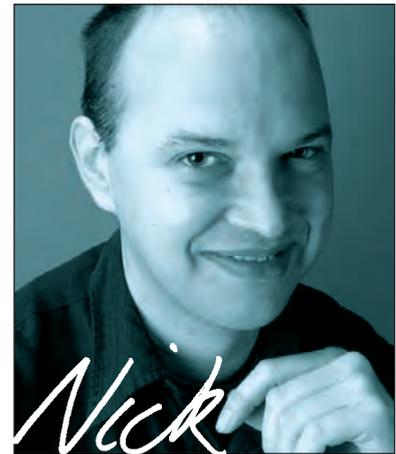
If you're more into the theory of the next generation of computers, you might like to check out our special What on Earth? feature this issue,

which tackles the seemingly impossible world of quantum computing and the physics of logic.

Or, if the thought of that sends cold electrons up your spine, you might like to check out the emulation feature which this issue tackles a range of popular consoles including the GameBoy Advance and the seminal PSone.

And though the ultimate roundup did consume 18 pages of the magazine, we still have plenty of space left for everything else. Three important new Linux distributions (Red Hat 7.3, Connectiva and Mandrake 8.2) go under the reviewer's magnifying glass, along with the latest version of Win4Lin, a server based Internet access control solution and a novel new motherboard from VIA.

On the tutorial front our regular programming series are still here, and are joined by the start of a series on PHP, as well as a guide to Gnumeric (though it should apply to any other spreadsheet). And don't forget to check out Linux Pro!



**Nick Veitch** EDITOR

# LINUX FORMAT

## Aims of the magazine

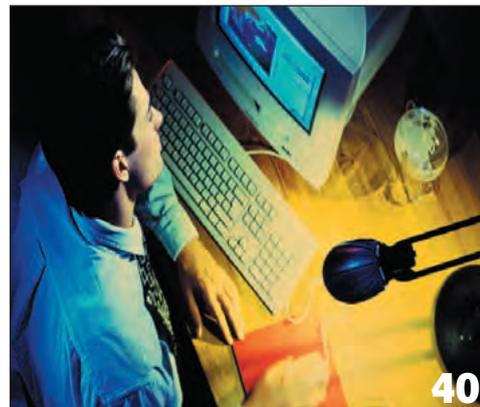
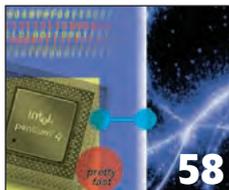
**Linux Format is a magazine dedicated to Linux and the Open Source community. We aim:**

- » 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.

**The biggest Linux office software roundup ever! p40**

**Quantum Computing – will physics reinvent computers? p58**

**Console cavalcade – a range of emulators for platforms as diverse as the GameBoy advance, and the PSone p62**



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



**Richard Smedley**  
Our resident leaf expert has an encyclopaedic knowledge of edible foliage. Handy.

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



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

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



**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.



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

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



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

**Maurice Kelly**  
Busy coder, electronic engineer and Midnight Oil fan – he still finds the time to read and write for us.

## 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

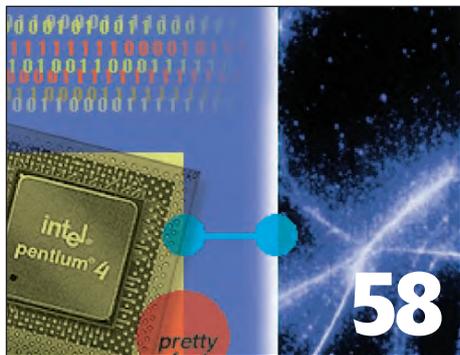
# LINUX FORMAT

LXF30 August 2002

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

## Quantum computing

Where particle physics and computing collide. A timely look at a subject that those working in the field claim not to understand.



## Games consoles

Console cavalcade – a range of emulators for platforms as diverse as the GameBoy Advance, Intellivision and the seminal PSone.



## Mastering Gnumeric

Our practical tutorial takes a look at using GNOME's spreadsheet, but you can apply the lessons to other Linux spreadsheets apps.



**COVER FEATURE**

# ULTIMATE OFFICE

Linux has been on the desktop for quite some time – read our ultimate office suites guide.

# 40

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

## »» REVIEWS

**18 RED HAT 7.3 PERSONAL**  
It's the corporate choice for servers – find out how the workstation-oriented version fares.

**22 WIN4LIN 4.0**  
More improvements from the Windows-on-Unix solution that is seeing competition at both ends of its market.

**24 MANDRAKE 8.2**  
The latest deluxe package of the newbie-friendly distro. We review the monster-sized boxed set.

**26 VIA MOTHERBOARD**  
A mini ITX, Eden powered motherboard, packed with goodies – how well does it perform under Linux?



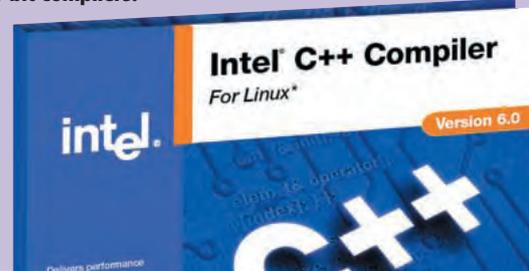
**28 SUPERSCOOT WEB FILTER**  
SurfControl port their advanced web filtering and caching software to Linux.

**30 CONNECTIVA**  
A distro which combines the ubiquity of RPM with the power of *apt-get*, and some config tools of its own.

**32 BOOKS**  
Help from O'Reilly for the stretched admin, the busy web programmer and the Qt developer.

## »» TUTORIALS

**68 INTEL C++**  
Speed up your C, C++ and Fortran code with Intel's 32 and 64-bit compilers.



**70 GNUMERIC**  
Our practical tutorial takes a look at using GNOME's spreadsheet to get some real work done.

**74 PERL**  
Using the popular Tk toolkit we help you to get graphical and add some widgets to your Perl programs.

**78 JAVA**  
Textual search and replace operations using regular expressions and patterns.

**80 KYLIX**  
This month we apply encapsulation, inheritance and polymorphism to our code as we introduce OO principles.

**86 PHP**  
Our new series starts you on the road to dynamic web content with forms and variables.

## »» REGULARS

NEWSDESK .....	6
MAILSERVER .....	12
HOT PICKS .....	35
ANSWERS .....	90
HELP WANTED .....	108
USER GROUPS .....	110

## »» FEATURES

HotPicks .....	35
From data mapping to war games	
What On Earth? .....	58
Is quantum computing?	
Emulators .....	62
Console cavalcade	
Help Wanted .....	108
Tux4Kids; GNOME; Rosegarden	

# Coverdiscs

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

Intel's C/C++ and Fortran compilers speedy x86 compilers; **GNOME** – rc2 of the free desktop environment; **123tkshop** flexible e-commerce solution; **Everybuddy** universal instant messaging; **Transcode** translates video formats; **Office suites** galore; **CUPS** printing made easy. *On the DVD:* **CPAN** The Comprehensive Perl Archive Network; **Gentoo Linux** the sysadmin's configurable Linux; **Boson** real time strategy.



# Newsdesk

United Linux; Red Hat and Oracle; Freecraft copylefted; Neverwinter Nights released; IBM gaming grid; nVIDIA gaming language; Mechanical CAD; Enterprise IDE; Walmart's Linux PCs; and OSS patents.

## UNIFORM BUSINESS DISTRO

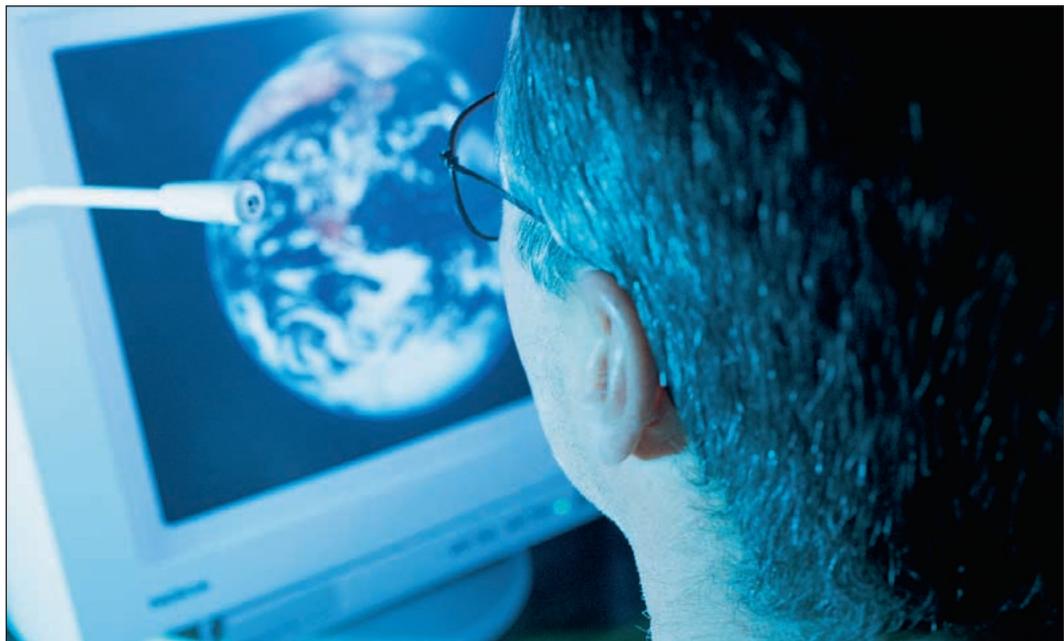
# United Linux takes to the world stage

**S**uSE, Caldera, Conectiva and TurboLinux have banded together to create a new standard base for business orientated Linux distros. The new über-distro, UnitedLinux, will hit the streets by the end of the year. All four companies have a strong toehold in their respective local markets but have struggled to make a real impact on the world stage, where 'Red Hat' is synonymous with 'Linux'.

This new initiative will, UnitedLinux hopes, "streamline Linux development and certification around a global, uniform distribution of Linux designed for business" and address the need for a standard business-focused Linux distribution. The four main players will contribute to the core operating environment which will then form the basis for their own branded, UnitedLinux certified, releases.

This release has not been without its controversies. The 'hand of friendship' was apparently extended to Red Hat and Mandrake, but was not taken and some critics said the whole idea was nothing more than a 'gang up on Red Hat' move.

Reaction to UnitedLinux was mixed – especially after a split emerged on the question of per-seat licensing – with Jon 'Maddog' Hall suggesting it may be the 'unique answer' to the problem of incompatible distribution. Richard Stallman released a call-to-arms "Licencing per seat' perverts the



GNU+Linux system into something that respects your freedom as much as Windows. Free software developers, please don't let them license your program per seat. Use the GNU GPL!"

The FSF added: "What they're doing is per-seat licensing, which is completely counter to freedom. They are adding on what they call value-adds, which is proprietary software, but we view it as removing freedom."

SuSE's US Director of Sales, Holger Dyroff, said that SuSE had no plans for a major shift in licensing policy.

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

## Unbreakable Linux

Oracle and Red Hat have become closer recently – they have cemented their relationship at a Larry Ellison hosted event in which the Oracle boss said his company would be 'recommending Red Hat': "We can't provide the same level of support for Oracle on other distros as they can on Red Hat," he said. He also announced 550 vendors have certified applications with *Oracle 9i* on Linux.

Unbreakable Linux is the integration of *Oracle9i Database Release 2* and Red Hat Linux Advanced Server and,

at the request of flustered IT managers, will break the "six month upgrade" cycle. Red Hat's Michael Evans said commercial software developers wanted a slower, more stable platform to work on, so Red Hat Advanced Server would have a 12-18 month development cycle and would receive extensive in-house testing at Oracle. But what of the cause of Ellison's Damascene conversion to Linux? Clustering, which Linux can bring into the hands of small/medium sized enterprises. [www.oracle.com](http://www.oracle.com) [www.redhat.com](http://www.redhat.com)



**Freecraft** – updated, and now protected by the GNU GPL.

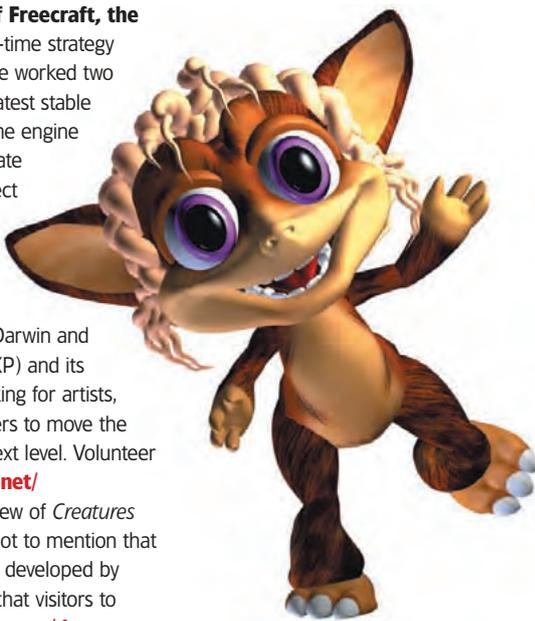
## STRATEGIC THINKING

# More games news

The developers of **Freecraft**, the cross-platform real-time strategy gaming engine, have worked two years to bring the latest stable version of their game engine (1.17) and to celebrate relicensed the project under the GPL.

The system currently runs on Linux, BSD, BeOS, MacOS/X, MacOS/Darwin and MS Windows (not XP) and its developers are looking for artists, musicians and coders to move the project on to the next level. Volunteer at <http://freecraft.net/>

Last month's review of *Creatures Internet Edition* forgot to mention that *Creatures* has been developed by Creature Labs and that visitors to [www.creaturelabs.com/ds](http://www.creaturelabs.com/ds) can download *Creatures Docking Station* for nothing.



Cute, evolutionary, and freely available with *Docking Station*.

## CROSS PLATFORM

# CA port application range to Linux

**Computer Associates** have begun porting almost their entire application range to GNU/Linux, a process that has seen them release 23 new applications this year, bringing their Linux application tally to 54.

Valerie O'Connel of research firm Aberdeen Group told Datamation that this change in focus was a big deal. "The result is an offering for Linux that is comparable to Computer Associates' capabilities in Windows 2000 and Solaris," she said.

## NEWSBYTES

■ The latest development kernels (from 2.5.14) have acquired Bluetooth support thanks to a project by Qualcomm which donated its protocol stack, *BlueZ*, to the open source community. The 'personal' wireless standard should be one of the headline features of the upcoming 2.6 release.



■ American thinktank the Alexis de Tocqueville Institution (AdTI) published, with great fanfare, a white paper called *Opening the Open Source Debate* which highlighted a wide range of problems and vulnerabilities in open source development. However it soon transpired that AdTI are financially supported by Microsoft and that the report rehashed much of the FUD surrounding the GPL that Steve Ballmer, Bill Gates *et al.* have been talking up over the past year.

■ Do you need a good stereo audio editor? *Audacity*, the multiplatform audio tool has finally reached version 1.0, which simultaneously hit the mirrors with the first beta of 1.0 which features higher quality sampling, more effects, more file format support and an improved UI. Both versions are available for Linux, Windows and Mac from <http://audacity.sourceforge.net/>

■ The follow up to withdrawn (but still available – it's GPL) video editing package *Broadcast2000* has reached Beta status, and users old and new are being encouraged to get involved in the testing of this new advanced compositing and editing tool. <http://heroinewarrior.com/cinelerra.php3>

■ Disney have become the latest in a long line of film studios to migrate its effects and animation shops to Linux. A deal has been done on hardware with Hewlett Packard. "Hollywood is at the leading edge of computing," HP's Martin Finkand said. "It shows what Linux can do."

■ The second generation GNOME desktop should be available as you read this. What was hoped to have been the final release candidate was on the mirrors as we went to press.

■ In response to an anticipated industry-wide move away from Solaris, Cadence Design Systems has announced plans to migrate all of its chip design tools to Linux by early 2003. Physical verification tools for 64-bit iterations of Linux should be available by June 2003, the company said.

## 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

# End of Linux is near

“As we are faced with the Redmond-esque monolith of United Linux, many feel that our beloved GNU/Linux is at the end of its life at the hands of Open Source poseurs. The glory days of the OS must be at an end; the suits have co-opted the future. That Linus himself touts proprietary software to administer the kernel source code is all the collateral needed.

However, the reality of the health of GNU/Linux is nowhere near as bleak. Developers, paid and unpaid, are advancing the state of the art at blinding speed in an evermore intimate dance with non-free code. Both have their place and both are learning to co-exist without extinguishing the identity of the other.

I was quite amused whilst installing nVIDIA video drivers since their choice is to provide non-Open Source drivers; my choice is to purchase their hardware and use their driver. As the install script completed, I was told that my Open Source system had been "contaminated" by the nVIDIA drivers. Perhaps in an ideal world it had been, but the two had cooperatively enabled a very nice video display for me. Likewise, the Apple-licensed code contained in the current release of *Crossover Office* provides me with crisply rendered fonts, something the free version of the library cannot do.

It is unlikely that I will ever be interested in using United Linux, but if some corporation wants to pay a per-seat license to use GNU/Linux with some proprietary applications, let them so choose; the community can always compile from source and learn to write their own. By cooperating, we are healthier than ever. ”

## ONLINE GAMING

# Linux gaming grid

**B**utterfly.net and IBM have deployed the first custom commercial grid for online gamers. The system was demonstrated at E3 in LA and is said to have taken two years of development using IBM e-business infrastructure technology. The Butterfly Grid intelligently redistributes computer power depending on the needs and popularity of a game or area, enabling Butterfly to support a vast number of players at any one time. Powered by a collection of rack-mounted IBM eServer xSeries systems networked fibre-optically and running Linux the IBM hosted grid can handle up to one million players on a 24/7 basis with an automatic fail over capability.

Scott Penberthy, IBM Global Services, said the system broke the mould, allowing efficient utilization of computing resources for high-performance 3D game-worlds. "We believe the Butterfly Grid is a breakthrough platform that will help entertainment, media and game companies reduce costs and better deploy their entertainment properties online."

[www.butterfly.net](http://www.butterfly.net)

COMPANY BUZZ TECHNOLOGY DOWNLOAD CONTACT

Butterfly.net provides the technology infrastructure for Massively Multiplayer Games (MMGs) that connect players on PCs, consoles and mobile devices. Our [Butterfly Grid Software Development Kit](#) is now available by download.

**NEWS**  
[Butterfly.net and IBM to Introduce First Computing Grid for Video Game Industry](#)

Butterfly.net is a production studio, online publisher and technology provider for persistent-state and session-based online games. Our first game is [Crvotids](#).

"Like everyone else, I was caught up in the games..."  
 John Keel, *The Mothman Prophecies*

Butterfly.net hope to host millions of players on their grid system.

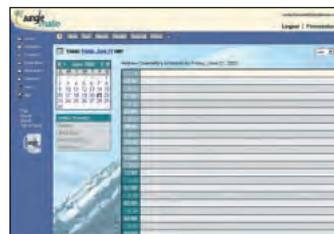
## Linux Web Watch/



Backflip – versatile bookmarks.



Yahoo! Briefcase – carries your files.



Junglemate – your online PDA.



GIMP – low bandwidth www.

## Your life online

Do you do all your surfing on one machine? No, me neither. So here's a collection of web spaces that will enhance your browsing from which ever machine (work, laptop, Internet café) that you happen to be using.

Bookmarks soon become essential if you spend any significant amount of time online, but their big drawback is they're so local. One of the more elegant solutions to this problem is [www.backflip.com](http://www.backflip.com) which lets you save URLs into a very Yahoo! like directory service.

You can sort your 'backflips' into various public and private folders and they can be accessed from any PC. You can even search through your collection. We had problems with

*Konqueror's* Javascript handling, but *Mozilla* and *Netscape* seem to work just fine.

When it comes to free online storage, we've still not found a better service than Yahoo's Briefcase ([www.yahoo.co.uk](http://www.yahoo.co.uk)), which gives you 30MB of space and an option to buy more. Its quite robust and is easily accessible from any connected PC. It too can have public and private folders making it ideal for sharing photos.

An online digital assistant can simplify access to addresses, diary and contacts, especially if you don't feel the need for a PDA. [www.junglemate.com](http://www.junglemate.com) provides all these services – and a few more – in a familiar looking web based interface.

There are hundreds of online calendar services available, but Junglemate's combination of task integration and clean interface make it a winner.

Finally, what would life be like

without leisure? If you're interested in some low-bandwidth fun, check out [www.guimp.com](http://www.guimp.com) which was created by the London Design Collective and claims to be the smallest website on the Earth.

It has both Flash and HTML areas and, in case you can't see the screenshot (yes, that is a screenshot up there, not a microdot featuring obfuscated Perl code :-), features a very faithful rendition of the game *Pong* (with sound!).

## NEWSBYTES



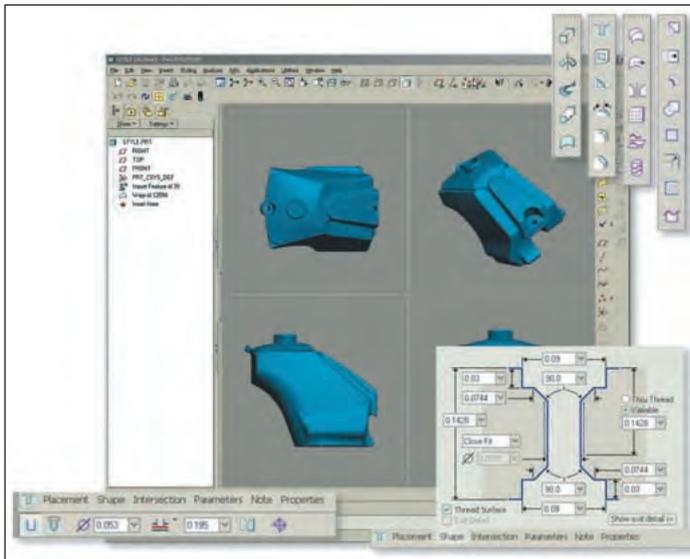
■ Mountain View Data are offering a free trial of their latest *MVD Powered NAS* (network attached storage) software *v1.5 Special Edition* using a streamlined Linux OS. MVD runs on Intel servers and supports clients running Linux, Windows and Mac OS.

■ NuSphere has developed 'the most advanced visual debugger for PHP available on the market today'. *PHPEd* is an IDE for handling PHP and a comprehensive set of editing, debugging and deployment tools.

■ The latest Netcraft survey shows PHP running 24 per cent of available websites, with growth predicted at 6.5 per cent each month.

■ The Linux version of the *Object-Oriented Graphics Rendering Engine (OGRE)* is 'almost here', according to programmer Thomas 'temas' Muldowney. *OGRE* is a "scene-oriented, flexible 3D engine written in C++ designed to make it easier and more intuitive for developers to produce games and demos utilising 3D hardware." Find out more at <http://ogre.sourceforge.net/>.

■ JD Edwards will be supporting Linux for the first time in the next release of its customer relationship management software package. While most ISV's have opted to join the Red Hat bandwagon, JD Edwards' first Linux version of CRM will be built on SuSE. However, the deal isn't exclusive. "If a customer is committed to Red Hat Linux, for example, we would work with them to support the platform of their choice," said VP of product management Lenley Hensarling.



**Pro/ENGINEER** – mechanical computer aided design from PTC and HP.

## MCAD FOR LINUX

## Coming soon: Pro/ENGINEER CAD

**PTC and Hewlett Packard** are joining forces to bridge another gap in the Linux application line up. *Pro/ENGINEER* is PTC's flagship mechanical computer aided design (MCAD) application which provides tools which aid the creation of "advanced, high-quality product models and designs that result in superior products."

PTC's Vice President of Technical

Marketing said the move to Linux was due to the growing demands from design engineers. "Adding Linux to the long list of certified platforms for *Pro/ENGINEER* ensures that our customers can take advantage of the innovations and competition taking place in the platform market to achieve the lowest total cost of ownership for their product development environment," he said.

## IDE FOR OPEN SOURCE DEVELOPMENT

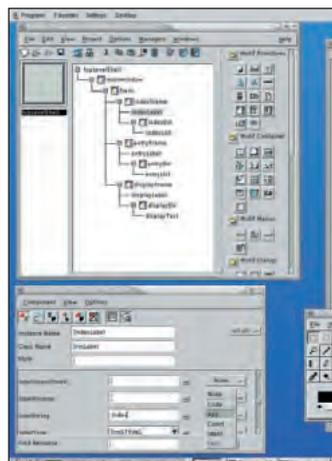
## Interface design made easy?

**Integrated Computer Solutions** have upgraded their flagship interface development environment *Builder Xcessory*. *BX PRO 6* is a 'major' upgrade, featuring a redesigned interface to make UI development easier, extensive support for a range of open source development tools such as *Autoconf*, *CVS*, and *Nedit* and the latest version of *ViewKit*, the integrated C++ object reuse framework with over 95 C++ components.

Peter Winston, president of ICS said that *BX PRO* users get the benefits and stability that long term development ensures. "Every year, there are new software

development tools that promise to change the world. Only a few like *BX PRO*, now in its sixth major release, have stood the test of time and continue to be extensively used in the development of enterprise applications," he said.

ICS is the market leader in *Motif* GUI development tools and are heavily involved in the *Motif* community, even sponsoring the *MotifZone* ([www.motifzone.net](http://www.motifzone.net)) which houses the public *CVS* source code for *Open Motif* and hosts projects for over 5000 active developers. You can find more details on their website: [www.ics.com](http://www.ics.com)



The latest version of **BX Pro** supports C++ object reuse.

## David Cartwright

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



## COMMENT Real IT education

“As I write this, the unis are chucking out their latest batch of graduates. As the years go by is that more and more students are able to afford their own computers and get connected to the Internet in their university residences – which actually lets them do stuff for real instead of having their practical experience limited to contrived examples in tutorials and lab classes.

As an employer I have seen graduates with the brains and the education but not the experience. They don't teach you how to set up a DNS at any university I know, nor what DHCP is, how to configure an *Apache* virtual server, and – most importantly – how technology really relates to business. It's reassuring, therefore, that more and more students are now able, thanks to falling technology prices and the vast amount of free, useful stuff for which we owe the FSF/Open Source movement a huge debt of gratitude, to do real computing at the same time that they're doing artificial and theoretical computing in their college courses.

I am increasingly able to employ new graduates without having to wonder whether they're going to hack it. I generally work with fast-moving organisations that rely on technology for their competitive advantage, and all too often I've found in the past that graduates know the theory but just can't "hit the ground running" in the real world. I know it's not much, but if you've actually used this stuff they teach you for real, rather than merely to answer the questions on a coursework worksheet, you're streets ahead of the others.

## ROLE PLAYING

# Never say never

**B**ioware have finally released *Neverwinter Nights*, the cross-platform RPG game that aims to be "the most complete software adaptation of the classic pen-and-paper Dungeons & Dragons role-playing game ever." Bioware's launch blurb claims the game should be good for between 60 and 100 hours of single-player gaming, but adds that it's more than just a game. "*Neverwinter Nights* breaks new ground as the first complete rôle-playing system for the computer platform, allowing its players to step into the shoes of the

Dungeon Master to tell their own stories and create their own adventures in the spirit of the exciting pen and paper original!"

Linux gamers will initially have to purchase the Windows version of the game to register at the community site and import essential elements when the native version becomes available on the *Neverwinter Nights* site. This should be soon after the official release. <http://nwn.bioware.com/>

Bioware have also released a dedicated Linux server for *Neverwinter Nights* <http://nwn.bioware.com/downloads/linuxserver.html>



Relive the late 70s D&D craze in glorious 3D technicolour.

## SUPERMARKET LINUX

## Wal-Mart's cut-price Linux PCs

The world's biggest retailer has announced a price drop in its OS-free PCs with the entry level machine now retailing at just US\$299. Users can also choose to have Lindows 1.1 preinstalled on the system and, as we go to press, Mandrake are in contractual talks to have their distro installed and 'running like clockwork' on the Microtel PCs. Microtel and Wal-Mart claimed that their decision to go with Lindows was based on the perception of ease-of-use and also its claim to be able to run the majority of

Windows software (a claim that has now been withdrawn from the Lindows website). However, a move to Mandrake could provide new users with a better range of support options, while retaining the user friendly image that Mandrake has cultivated.

Wal-Mart caused a storm when it became the first – and biggest – retailer to offer customers something other than a Windows PC, and analysts are predicting that other vendors will be watching the experiment with interest.

**Microtel PCs with LindowsOS**

The low cost alternative to computers preloaded with Microsoft Windows. These PCs ship with an exciting new Linux based Operating System (OS) named Lindows. This exciting new OS delivers the stability of Linux with the ease of Windows and they include a trial membership to a library of over 1,000 software programs so they can be outfitted for any purpose whether business, home, or entertainment.

Sort this list by: Top Sellers

Items 1-8 of 8 Total

 <p><b>NEW Item!</b> \$299.00 <a href="#">Add to Cart</a></p>	 <p><b>NEW Item!</b> \$599.00 <a href="#">Add to Cart</a></p>	 <p><b>NEW Item!</b> \$499.00 <a href="#">Add to Cart</a></p>
--	--	--

Lindows finally gets beyond its vapourware tag. But is it about to be replaced with Mandrake?

## Embedded Linux News



● *Mozilla* has been garnering all the headlines lately, but it's not the only game in town, and if you have older hardware or just want something that won't hog your system, *Dillo* may be the answer. This diminutive browser was built for handhelds but will compile and run on just about any flavour of Unix you care to try. Written completely in C, the source comes in at 300k and the binaries at 200k and the whole thing has been released under the GPL. <http://dillo.cipsga.org.br/>

● Austria's BMS Bayer have developed a portable device for monitoring levels of X-ray exposure in medical environments; the data can then be processed and stored to comply with recently enacted laws in the US and Europe. The EasyDose device is built around a Cyrix GXM processor and uses Transmeta's MidoriLinux Distribution housed on a 32MB CompactFlash card. Norbert Bayer, the company's CEO said using Linux enabled the company to concentrate on the application "without having to carry

an enormous unnecessary OS payload"

● If a Mercedes just isn't enough, DaimlerChrysler have shown off a Merc with an integrated Linux system that powers two 15 inch TFT screens for back seat passengers and a central console for the driver. Applications demonstrated in the car included MP3 playback, radio, navigation, Internet access and maintenance services. Project leader Jentro claims this is the first car equipped with Universal Mobile Telecommunications System (UMTS) services. [www.jentro.com](http://www.jentro.com)

● Finally, if the Merc is too expensive, the latest Linux handheld, the WalkPad GP1288 is priced at the equivalent of \$160.. It sports a 58MHz MIPS processor, 32MB of Flash ROM and 16MB RAM. [www.ggv.com.cn/product/indexWP.php](http://www.ggv.com.cn/product/indexWP.php) In English at: [http://babelfish.altavista.com/tr?hidden=done&tt=url&url=http://www.ggv.com.cn/product/indexWP.php&lp=zh\\_en&Translate](http://babelfish.altavista.com/tr?hidden=done&tt=url&url=http://www.ggv.com.cn/product/indexWP.php&lp=zh_en&Translate)



**BIG BUCKS LINUX**

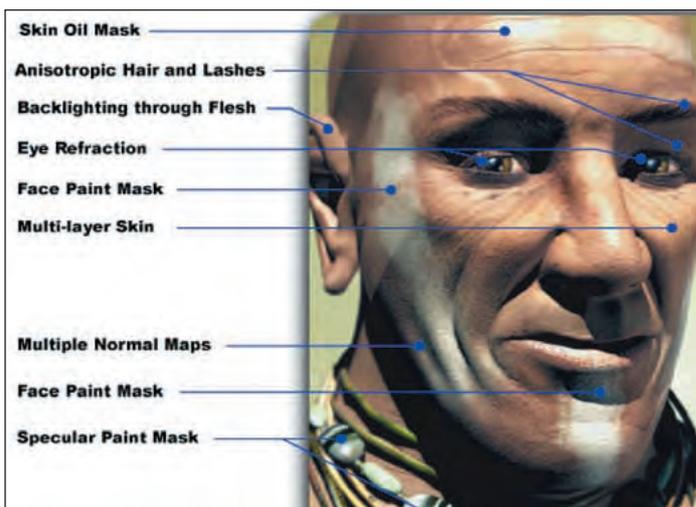
# What IBM giveth...

**Two recent announcements from IBM** highlight contrasting success for the company's Linux strategy. The first was the abandonment of its initiative to push Linux as an OS option on a range of Thinkpad brand laptops, for which IBM cited lack of interest. Hand in hand with this announcement was the launch of a New York-based centre which will provide solutions,

training and education on Linux and open source software for the financial services sector. Big Blue hope the new centre will lure corporate users away from Sun, who are traditionally strong in this sector, and onto Linux and its range of eServers.

IBM will sink \$1 million into the project and, with partners including Sybase, Veritas and SunGuard, will

provide a means for enterprises to evaluate, test and deploy a range of Linux apps. IBM said financial services were migrating to Linux on an unprecedented scale, citing the examples of Merrill Lynch, UBS Warburg, Lehman Brothers, E\*Trade and Credit Suisse First Boston who were all actively porting systems to the operating system.



**Cg** will make the creation and manipulation of complex 3D scenes easier.

**GRAPHICAL LANGUAGE**

# A revolution in games creation?

**nVIDIA has developed a high level programming language** that could lead to an explosion of cross-platform 3D games. Dubbed Cg ('C' for graphics), the language has been specifically designed to handle the creation and manipulation of 3D graphics, and comes with its own toolkit consisting of a Cg compiler, browser, standard library and a variety of pre-written shaders. The compiler can output with DirectX or *OpenGL*, is supported on Linux, Windows, Macintosh and Xbox and has been partially open sourced.

As well as being backward compatible, nVIDIA claim the compiler will support next generation graphics chips, meaning code can be recompiled to take advantage of future features without much intervention

from a programmer. "Cg is a significant milestone for the computer graphics industry as it moves real-time 3D graphics programming a quantum leap closer to the programming model of cinematic rendering," said Jen Hsun Huang, CEO at nVIDIA. "It will do for GPUs what C and C++ did for CPUs."

Over 100 games developers, including Blizzard, Sony and Valve, have so far pledged their support for the standard, as have application vendors such as Alias|Wavefront. Gabe Newell, from Valve Entertainment (developers of Half Life) said the language would revolutionise the way games are created.

A public beta of the Cg toolkit for Linux is available at <http://developer.nvidia.com/>

## PATENT CONTROVERSY

# Good? Bad? or merely necessary?

**Risking the wrath of the Open Source community** Red Hat have applied for a pair of software patents in what the company called a defensive move to protect the Free Software community. In a statement, the company said it has consistently taken the position that software patents generally impede innovation in software development.

But "we are forced to live in the world as it is, and that world currently permits software patents. A relatively small number of very large companies have amassed large numbers of software patents. We believe such massive software patent portfolios are ripe for misuse because of the questionable nature of many software patents generally and because of the high cost of patent litigation." The best defence against such misuse, Red Hat claim, is to develop a corresponding portfolio of software patents for defensive purposes. "Prudence dictates this position."

Patents, covering 'Embedded Protocol Objects' (20020049834) and 'the method and apparatus for atomic file lookup' (20020059330) are designed to prevent proprietary developers from taking the freely available code and patenting it.

[www.redhat.com/legal/patent\\_policy.html](http://www.redhat.com/legal/patent_policy.html)

**Jono Bacon**

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



## COMMENT

# The web revolution

Technology is something that has always been able to fool and astound the masses, and this humble student is no exception. In the past few years we have seen a great change in technology as the computer, the interface, the global economy and the Free Software world has been thrown in a bag to settle their differences. The question I have been thinking of recently is whether all the bickering over Operating Systems, interfaces and the economy is worth it; we are a web-enabled world after all.

Explosive growth in popularity of the web has naturally led to the development of a number of technologies to solve increasingly complex problems. PHP, Python, JavaScript, DHTML, XML, XSLT and XMLRPC are all extending areas in ways we never expected. Add to this high speed and broadband Internet access and we have a powerful infrastructure.

The future looks bright for the web, but may actually cause an implosion for the use of a PC. In future years the PC may become a mere terminal in which an amazingly capable browser will run everything. Imagine a browser with a built in *OpenGL* library for 3D game sites to run online games, imagine making music online with a browser capable of running a sequencer and imagine your hard disk being a server across the world. Is this the way computing is going in the networked age? Will the browser become the analogous OS and Linux become the BIOS?

I think this vision is not totally opaque and may present us with not only technical challenges, but also the social challenges of the networked world.



# Mailservver

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)

## Burning issue

Firstly, thanks for a great magazine! As long as the content continues to be as informative and challenging; I will continue to subscribe.

I read frequently that there are questions regarding the creation of CDs from the ISO images provided on the DVDs supplied with *Linux Format*. So I consider myself fortunate, as I haven't encountered any problems using any of the images from the DVDs. I also notice that most of the readers appear to use *Nero* to burn the CD from within Windows.

I thought that I would share some of the things that I do to create the CDs, in the hope that it might help, as I too use *Nero*. My recipe for burning ISO images to CDs from within Windows is:

- 1 Shut down as many or all utilities in the *Systray* (next to the clock) as you don't want any interruptions (virus scans, connections to get email, ... whatever!).
- 2 Copy the ISO image from the DVD to the hard drive. DVD (and CDROM) drives take a while to come up to speed, and the speed varies as they read data, while the hard drive runs all of the time. So eliminate the possibility of delays in the data transfer.
- 3 Start *Nero* and cancel the Wizard that usually pops up.
- 4 Select "Burn Image ..." from within the "File" menu, and choose



## ★ Letter of the month

This month's winner receives a copy of **Red Hat 7.3 Personal**

### First Computer

Hi, I am working on the Colossus Rebuild Project at Bletchley Park and I would like to put the record straight.

There have been several letters about the article on Encryption (*LXF21 Dec 2001*) and Enigma, I would just like to put the record straight.

Before the second war the Polish had indeed cracked the Enigma code, but when the Germans increased the number of permutation on the Enigma, the Poles were no longer able to decipher the German Enigma Codes. At the beginning of the war Bletchley Park (Station X) was given access to the techniques used by the Poles in deciphering Enigma, and in 1940 the Bombes were developed to aid the decryption process. These Bombes were Electro-mechanical devices that were definitely not computers; they were basically motorised switches and were not capable of decrypting the Enigma messages. Each Bombe was basically simulating 12 Enigmas and each bank of switches represented a single Enigma machine. The Bombe looked for a mapping between 3 letters (switch setting), when the mapping was found, the switch setting represents the wheel setting

on the Enigma, with this information the message can be easily be decoded.

The German High Command's telex messages used the Lorenz teletypewriter-based cipher system; it was a rotor-based system like Enigma but far more sophisticated and was a primitive digital device. To help with decoding these messages Bletchley Park had a device known as the Heath Robinson, which was electromechanical with 2 paper tapes loops, but it was not very successful due it being slow, and problems with synchronising the 2 paper tape loops. Now Tommy Flowers, a GPO development engineer at Dollis Hill, heard of the problems with the Heath Robinson, and he set about the development of an electronic device to replace it. In 1943 at Dollis Hill Colossus was born, the first electronic computer, it was shipped to Bletchley Park for evaluation, and it was so successful that 11 more were built at Bletchley. Colossus had 2,500 valves and consumed 7kw of power when running and was made using standard GPO telephone exchange components, it had what we know today as logic gates and it was capable of parallel processing.

When the war ended, 10 of the 12 Colossuses were deliberately smashed to smithereens and all



the documentation was burnt, the remaining 2 ended their life at GCHQ and were also destroyed in the 1960s. Everyone at Bletchley Park was sworn to keep Bletchley Park a secret and only recently has it come to light as to the ground breaking work that went on there.

So the first electronic computer was developed by Tommy Flowers in the UK, and was used as an aid in the decryption of the Lorenz telex.

Bletchley Park is now a living memorial to those outstanding achievements, and as visitor to Bletchley Park you can see Colossus at the end of the guided tour.

For more information [www.bletchleypark.org.uk](http://www.bletchleypark.org.uk)  
Phil Hayes (*Colossus Rebuild Project*)  
Thanks for all this information, and the website link which is definitely worth a look. I'm a bit confused now, did it turn out that our original piece was right then?

Anyway, for taking the time to set out all this excellent detail for us, please accept our prize this month – a boxed personal edition of Red Hat 7.3

the ISO image stored on the hard disk that was copied previously.  
5 *Nero* recognises the format as "foreign", and in the newer versions there are some options "greyed out" in the Dialog Box that pops up.

Beware that earlier versions of *Nero* allow you to change the "Mode" and "Block Size" which results in a "problematic" CD, or "coaster," if you choose to change them. Stick with the defaults, usually Mode 1

and a 2k Block. If it's "greyed out", well, you can't change it anyway.  
6 Check again the speed of the CDR to be written. No point in setting 24x write if an older CDR was made to be written at a



maximum of 8x, you will get errors. I usually write the CD at 8x, which only takes 10 minutes, as the bulk pack of CDRs I have will not support more than 8x write.

**7** Check the box to "Finalise" the CD (no further writing to the CD is possible) or it will not boot or be read on older equipment, regardless of the age of the BIOS.

**8** Write the CD.

**9** If all went well, Nero should report that the CD was written successfully. You should be able to view the contents of the CD in Windows Explorer, and it should boot (if it is a bootable CD, and the BIOS is set to boot from CD).

I appreciate that everyone probably creates the CD in the manner that I have described, but I've been caught out by the CDR speed and by "messing" with the "Mode" etc., so it's worth checking those items. Now for my wish list! What about a review and update for the Slackware Linux distribution?

Ellis Skinazi, *via email*

Many thanks to yourself, Roy, Stuart and others who also mailed in similar instructions. I hope those Nero users are taking it in.

As for Slackware, there hasn't been a significant release since June of last year, but I'm sure we'll be reviewing 8.1 when it is finalised.

## DVD ideas

Although I have read your magazine for quite a while it is only recently that I took out a subscription. What convinced me to do so were the DVDs that were now accompanying the magazine. It contained software that would have been a pain to download. Anyway, if you are ever looking for ideas of what next to put onto the DVDs here are a few suggestions:

### a. A copy of CPAN

This area is such a great resource for people "playing" with Perl. I am constantly pulling down libraries from CPAN. CPAN on DVD (or CD) would be a good toolbox to own.

### b. A copy of the World Of Spectrum website.

Due to nostalgia I keep an eye on this website. Occasionally I download the odd snapshot. However, I believe, the maintainers have limited the download number to keep (understandably) the



We welcome your suggestions for DVD content.

bandwidth on the site down. Having the contents on DVD would be one way of reducing this load.

**c. A copy of one of the mirror web sites where GNU software is contained.**

Often I pull things down from a website that contains GNU software. The software itself may be destined for Linux itself or another Unix like box. Some of these websites are slow (as well as my Internet connection). Having such software on CD/DVD would make it much easier to access.

**d. A copy of a web site where Xemacs is.**

Xemacs is such an important part when I am developing code. Many Unix installations do not come directly with Xemacs and those that do usually don't include files from the contrib directory. Having such contents on disc means that I can reach for it when it is needed.

Jon, *Australia*

Thanks for the suggestions.

a) Done – take a look at the DVD this issue

b) Hmm. I'll take a look, but the Linux interest here is?

c) We have done this partially in the distant past, and it can be very large, especially if you get the documentation too! Are there some specific programs you're interested in

d) Well, we could add this to our essentials directory – but does everyone else think it's essential?

## Price check

I ordered the *Linux Cookbook* recently reviewed. The price turned out to be £29.99, not £23.99.

Eric Richards, *via email*

Price checker Richard Smedley replies: Yes, the publishers didn't have a UK RRP, so I put in a market price. I have seen it cheaper than that, though of course one must take into account postage cost and the limited periods of some offers from online retailers. Thanks for your comment.

## Availability

I am a recent addition to your reader base, and very much enjoyed the latest (May) edition.

The experience would have been much more rewarding if I had been able to find a copy of the DVD edition, however. I've spent hours driving around Central Scotland, but in the end had to settle for the CDROM edition – most frustrating.

Are you certain you have the ratio of CDROM to DVD versions you issue correct? I don't want to subscribe to LXF just yet, but it would've been nice to be able to buy the DVD edition in the interim! Keep up the good work.

Martin, *Scotland*

The numbers we print and sell of both CD and DVD versions are closely monitored every month and adjusted accordingly. Unsurprisingly, the DVD edition becomes more and more popular as time goes on, so we do

**"In 1943 at Dollis Hill, Colossus was born, the first electronic computer – developed by Tommy Flowers, in the UK, to decrypt the Lorenz telex"**

print far more of these now than we ever did, though there are occasional shortfalls when we produce a particularly popular issue, like issue 27. I'm sorry you had trouble obtaining a copy of this issue. If you have problems in the future, you could always order a back issue from our warehouse – email [backissues@futurenet.co.uk](mailto:backissues@futurenet.co.uk)

## Short and...

I really liked the security roundup in LXF27. It was the best part of the issue IMHO. Hopefully we'll see security and security tools featured regularly in *Linux Format*. BTW, great mag!

Mark Hill, *via email*

We're glad you liked it. Take a look at *Linux Pro* for features and tutorials on security issues. We should also be covering firewalls in the main magazine soon.

## Last word on glue

I'm somewhat surprised that so many people complain about the glue used to join the CDs to the magazine. If you put the CD cover into the fridge for enough time and then you push gently with your thumb along the strip it pulls off in a nice little ball with no effort and no traces. It worked nicely until the Italian distributor unfortunately chose to use a sticky tape instead, which is worse.

Rodolfo Cerreto, *Italy*

The title says it all.

## Open portals

I would be interested if you are considering running an article on OpenSource CMS and Portals. Seeing as *Linux Format's* own site uses one I think it would be worthy.

These solutions are forming a major part in being able to provide professional looking websites to those without the skills to do so. The ones I would consider for review are:

Mambo SiteServer – [www.mamboserver.com](http://www.mamboserver.com) and [www.miro.com.au](http://www.miro.com.au)

PHPNuke – [www.phpnuke.org](http://www.phpnuke.org)  
PostNuke – [www.postnuke.com](http://www.postnuke.com)  
ezPublish – [developer.ez.no](http://developer.ez.no)

I personally use Mambo and find it the easiest to install, maintain and customise. I would be interested to read your views. I look forward to your response.

P.S. Love the *Linux Pro* 'mini mag'.

Robert Castley, *via email*

This is an article idea that has been on the backburner for a while at *Linux Format*, but keeps getting put off for one reason or another. It would certainly make for interesting reading.



As you notice, the LXF site itself runs on *PostNuke* currently. We did run on *PHPNuke* before that, but the 5.3 release took so long to come out, we switched to *PostNuke* so we could have some working forums!

Can anyone else make any recommendations?

## CR/LF

I hate to nit-pick, but I've just finished reading the June edition of *Linux Format* and found two otherwise very interesting articles (both by Charlie Stross) marred by a basic inaccuracy.

In both the Mac Interoperability article (page 58) and the Perl tutorial (page 75) he claims that Unix files use a single carriage return (described in one article as ASCII 10 and the other as ASCII 13), as an end-of-line marker for text files whereas DOS/Windows files use a carriage return/linefeed pair.

Every Unix system I have worked on, from the mainstream Solaris to the idiosyncratic version on the old pre-HP Apollo workstations (yes, that does date me somewhat!), including Linux, uses a single LINEFEED character (ASCII 10, hex 0A, format code `\n`) as an end-of-line marker for text files. DOS/Windows do indeed use a carriage return/linefeed pair (in that order).

I'm afraid I can't comment on how Macs treat text files as they aren't common in my industry sector and I haven't had much contact with them. The `to_dos` and `from_dos` Perl functions on page 75 will, in fact, do the right thing, so I hope that the accompanying text was just a typo/printing error, rather than a basic misunderstanding...

Nigel Collins, *via email*



### CR/LF – Linux uses one, Macs use the other and DOS uses both!

Thanks for the correction I think we must have got our keyboard interface wires crossed. The Mac does use a single carriage return, just to be different.

## When will it end?

I'm getting bored. I'm bored of people writing into magazines and websites to inform us that Linux is not ready for the desktop, and telling us what Linux needs to do to make it on the desktop.

I'm prompted to write this email based on your latest "Letter Of The Month" in issue 28, entitled *Linux GTI?* by Chris White. I'm not arguing with what Chris is saying – he makes some decent points. I just think it's time we stopped bleating on about how LINRFTD (Linux Is Not Ready For The Desktop). It's a well-known state of affairs, and it's beyond time that people stopped talking about it, and did something constructive.

I'm very disappointed that this old and tired opinion was awarded "Letter Of The Month." It contained nothing new, and is not even the best opinion I've ever seen on

the subject. And you decided to award it, instead of Mika Laaksonen's excellent and informative piece on code-breaking and Enigma machines.

Maurice R. Kelly, *Belfast*

Possibly not, but we would like to encourage people to write in, even if many would not share their opinions. I hope you approve of this month's winner. And you'd better close your eyes for the next bit.

## Flame war return

It seems rather ironic that in an issue where you have a column *Hooligans – Time To Stop*, by Hoyt Duff, you then proceed to do the same to Steve Townsley – *A big Question?* Is this the normal defensiveness that I come across regularly with Linux fans (I consider myself one and have a dual boot Windows/Linux PC with both operating systems being used probably 50-50) whenever someone criticises an aspect of Linux, especially if it is in comparing it with an aspect of Windows. It doesn't really matter why someone wants to try out a piece of

software, after all how can you compare differing features without actually trying them out and yes installing a piece of Linux software is often a lottery.

The 'question' in the letter is the very same one I ask myself all the time when installing a piece of Linux software. Surely it is not beyond the wit of someone in the Linux community to get everyone to agree on a standard installation method. Until this happens many who would consider Linux will be dissuaded by the extra work often involved in installing a piece of Linux software compared to installing a piece of Windows software (whoops! does that mean that I will get a snotty reply from your letter pages editor).

Until both the installation of Linux itself and other pieces of software you may require becomes as easy as it is with Windows, Linux will sadly remain an also ran. Linux is moving in that direction but it is a long way off yet and criticising the messenger is not the way to go, grow up.

John Phillips, *via email*

## Steve Townsley

Having been mentioned three times in the letters page of the July issue I thought I would give my response back.

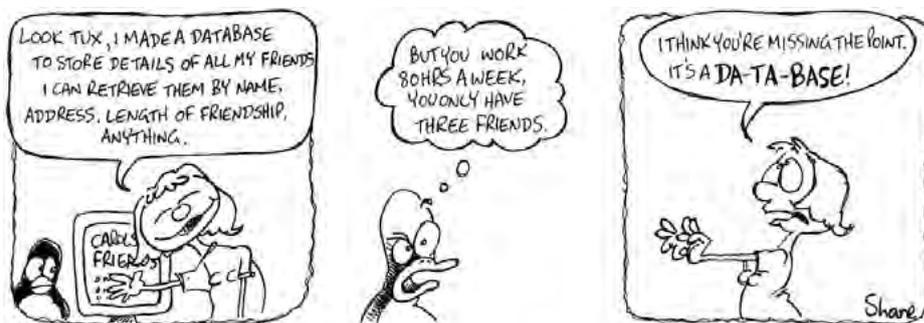
Firstly, like other correspondents, I thought your original reply to my letter was flippant and arrogant. Secondly, my main point was not about *qmail* or any other particular program. My main point was in fact the importance of standardising Linux to the extent that programs run on all distributions.

This is not some newbie opinion expressed as a non-technical reader. Unix failed to dominate

# Helpdex

BY SHANE COLLINGE

shane\_collinge@yahoo.com





outside high end enterprise servers in 70s and 80s because it split into multiple flavours. Windows has succeeded (in part) because it has unified binaries that have a well understood installation and de-installation method. Thus the average desktop computer user is not using a version of Unix but a version of Windows.

There are many things wrong with Windows. However Microsoft has made attracting a mass market a priority. Linux hasn't. I don't say that in an insulting way. It simply wasn't an objective of Linux when it was written. I first saw Linux in 1994 when you practically had to build a boot floppy from DOS to run it. Today's users expect to download a program from the Net and for it install first time and

**“In 1994 you practically had to build a boot floppy from DOS to run it. Today's users expect to download a Linux app from the Net, and for it to just work”**

work. They are not prepared to build programs from source code or worry about what library program “X” needs.

In my original letter I used a struggle with *qmail* as an example. I could have chosen a hundred other products. You dismissed the exercise by saying “oh you should have just bought a commercial email program.” Yet not only do you promote hundreds of programs on coverdiscs for self-install by readers, you also added an ironical touch to your comment by printing an article in the same issue as my original comments, that was

describing attempts to get a common Linux standard.

I have fun with Linux. I am a computer hobbyist of many years and I work in the IT industry. I am saying that a united Linux binary format with a consistent directory structure and a consistent install procedure would benefit Linux users of all abilities and might even bring Linux to the home user desktop.

I might even say that as a magazine publisher, such a product could bring you more readers. You, a magazine supporting Linux users, should be leading the call for a Linux that more people can more easily use. Instead I got the comment that normal people shouldn't spend time installing *qmail*, they should buy

a commercial product.

If Linus had taken your advice he would have gone down to a software store and bought Windows 3.0 rather than write Linux at all.

Name, *via email*

My apologies, Steve, if I was rude to you. To be honest, this could have been any number of letters that we receive regularly on the subject of installing this or that. I'd also like to point out that I personally do agree with the idea that software should be easier to install, and the magazine does cover ideas like LSB as you noted.

Where I think we disagree is that the problems people have installing software are inherently something to do with Linux. This simply isn't the case. It is eminently possible to write click and go installers for software, as evidenced by Loki's game titles, Sun's *StarOffice*, Codeweavers *Wine*, *Mozilla* and plenty of other software. Lack of uniformity in directory structures may complicate the matter, but it doesn't prevent installers from working.

The problems you describe are actually more to do with Free Software (because it almost always has to be compiled) rather than Linux, and they are the same whether you are using Windows, Linux or whatever – just try installing the latest *Emacs* source on a Windows box.

The magazine does provide plenty of help, tutorials and advice on compiling software, using package managers and anything else we can think of, but yes, it is still a problem, and not an easy one to solve. If you want to argue that Free Software doesn't have a future on the desktop because it's too hard to install, or too confusing then you have a better argument. [LXF](#)

### Submission advice

#### WHAT WE WANT:

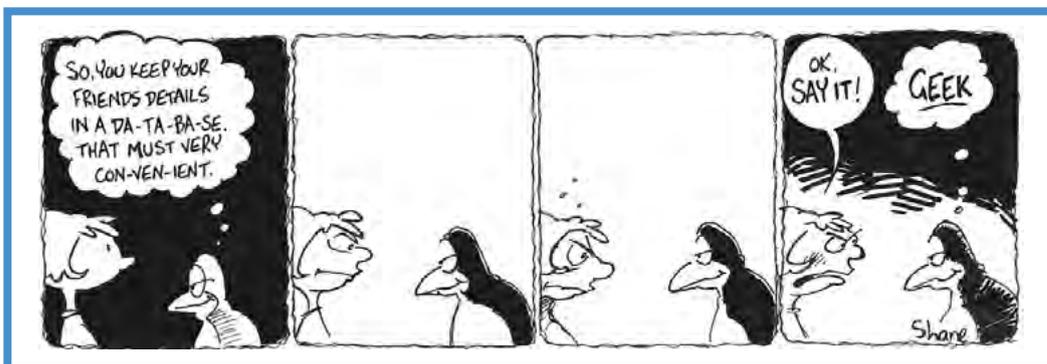
- Letters about the magazine or Linux in general
- Constructive criticism
- Your opinions
- Concise points about relevant subjects

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



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

SuSE Linux Business Solutions

# ALL YOU NEED

FOR INFRASTRUCTURE

SUSE LINUX ENTERPRISE SERVER



Security • Stability • Scalability  
Efficiency • Reliability • Flexibility

A professional cross-platform Server environment that is guaranteed to offer a long-term robust and reliable system for your significant business applications.

- cross platform certification
- advanced maintenance programme
- proven enterprise implementations
- world-wide support available
- planning and investment security through extended version support.

### WANT TO KNOW MORE?

CALL YOUR SuSE TEAM NOW AND ASK THEM HOW THEY CAN HELP YOU!

**020 838 74088**

INQUIRIES FROM VARs AND ISVs WELCOME

SuSE Linux Ltd.

The Kinetic Centre ·  
Theobald St.  
Borehamwood  
Herts. WD6 4PJ  
Phone 020 838 74088  
Fax 020 838 74010  
[info@suse.co.uk](mailto:info@suse.co.uk)  
[www.suse.co.uk/sles](http://www.suse.co.uk/sles)



BUSINESS  
PARTNER

# 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...

### Red Hat 7.3 Personal >>

It's the corporate choice for servers — find out how the workstation-oriented version compares **p18**

### Win4Lin 4.0

More improvements from the Windows-on-Unix solution that is seeing competition at both ends of its market **p22**

### Mandrake 8.2 ProSuite Edition

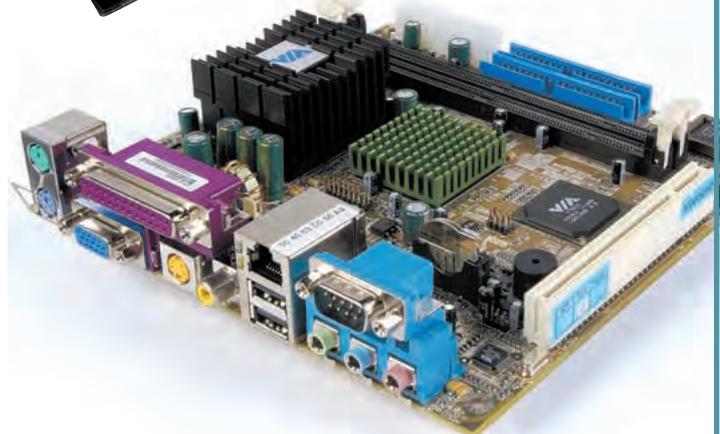
Does the latest huge box of tricks from Mandrake enable them to keep their 'Newbie-friendly' crown? **p24**

### Epia Eden 5000 >>

A tiny, low-consumption, fanless mainboard with many uses — how well does it perform under Linux? **p26**

### Superscout Webfilter

SurfControl port their advanced web filtering and caching software to Linux. Has it made the leap successfully? **p28**



### Conectiva Linux 8.0

A distro which combines the ubiquity of RPM with the power of *apt-get*, and adds in a few config tools of its own **p30**

### Books

Help from O'Reilly for the stretched admin and the busy web programmer, as well as Qt development **p32**

## COMING UP SOON...

### Zend Studio

PHP developers deserve an IDE too, so we'll be looking at this professional environment.

### Quasar Accounts

So you think there are no accounting packages for Linux? Then think again.

### ELX Linux

A new distro, aimed at stealing the ease-of-use crown from Lycoris. We'll be reviewing it soon.

### Slackware 8.1

Does the old man of Linux distros have anything to offer in their new release? Watch this space.

### JRun 4

Expect a review of Macromedia's acclaimed Java application server in an upcoming issue.

### VMware ESX Server

It's like VMware Workstation, but for servers. Can it solve all your management issues?

## WORKSTATION DISTRO

# Red Hat Linux 7.3 Personal

**Richard Drummond** asks whether the biggest is necessarily the best as he examines the latest release from Red Hat.

Workstation-oriented box of the popular server distro. Desktop rivals are SuSE and Mandrake.

- **PUBLISHER** Red Hat
- **WEB** [www.redhat.com](http://www.redhat.com)
- **PRICE** £49.99

For many users Red Hat is synonymous with Linux, so the quality of a new Red Hat distro will pretty much determine how Linux is judged in the corporate world. Thankfully, despite the odd mistake, Red Hat have built a solid reputation by supplying robust and well-tested distros and good support, and Red Hat Linux 7.3 provides more of the same. You won't find much innovation here, and, as far as the desktop goes, rivals such as Mandrake and SuSE offer a better-integrated and more comfortable user experience; but if you want dependability and the assurance of compatibility that using Linux's most popular brand brings, then Red Hat is an obvious choice.

Red Hat 7.3 ships in *Personal* and

*Professional* editions – although, there's actually little real difference between the two. The former is being reviewed here, and this is ostensibly aimed at the workstation user, whoever that might be, while the latter is targeted at the server market. Both editions are in fact the same core distro, which is provided on seven CDs (three install discs, three of source code and one of documentation) and

this is the same as the ISO images that you can download from the web for free. Of course you get some extras in the Red Hat *Personal* box: a 30-day membership of the Red Hat Network, an installation manual, a *StarOffice 5.2* install disc and a bonus CD of proprietary applications (this is mostly evaluation software, however, and the only thing of value to most here is Sun's JDK). Similarly, the *Professional* edition has an equally superfluous disc of server applications. You also get a DVD which you can install from without all that tedious disc swapping.

## Put on your Red Hat

Red Hat don't make a lot of noise about ease-of-use, but Red Hat 7.3 is relatively painless to install. As well as offering a traditional install, the first install disc can be used to upgrade an existing Red Hat install, do network installs via NFS and also serves as a rescue disc. Since the release of Red Hat 7.0, you've been offered a choice of graphical or text-based installation, and the GUI-based installer has matured considerably. In 7.3, this is visually appealing and mostly straightforward – although navigation, online help and keyboard short-cuts could all do with being improved. Like the rest of Red Hat Linux, the installer gets the job done with the minimum of fuss and few embellishments – but Red Hat would benefit from learning a

few lessons from their rival Mandrake.

For example, Red Hat's partition editor, *Disk Druid*, still offers no tools for non-destructively resizing partitions, despite being based on Parted which does support this – my guess is that this omission is a user support issue, rather than a technical one. Otherwise, disk management is quite flexible. Support for creating *ext2* and *ext3* filesystems (but no *ReiserFS*, alas), software RAID and LVM are included, and the Red Hat kernel has been patched to cope with disks larger than 137 GB. When using the installer to upgrade an existing system to Red Hat 7.3, you can even migrate partitions from *ext2* to *ext3*.

Hardware detection is equally solid. My only real complaint is that I found that the graphics card configuration could easily become confused by multi-headed systems. On the whole driver support is good – especially for networking and SCSI hardware – but once again Red Hat is conservative. Their strict open source policy for the core distro is laudable, but this means that there are no closed or half-open source drivers included – such as NVidia's 3D-accelerated X drivers or support for Lucent and Conexant 'winmodems' – like some distributions are beginning to offer.

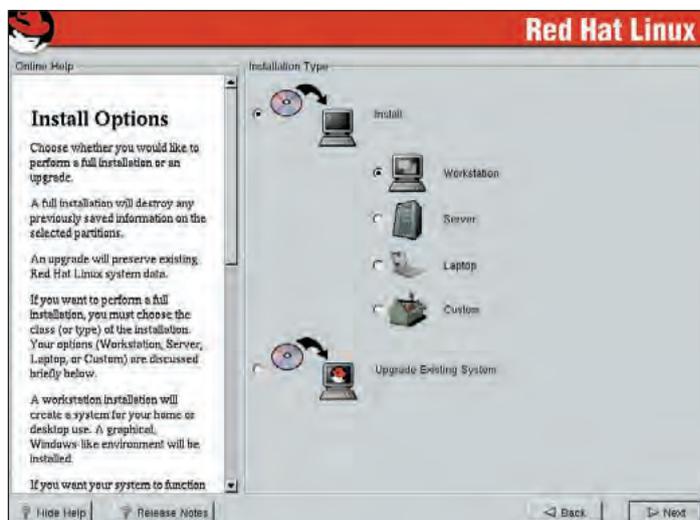
Like all RPM-based distros, choosing which software packages to install can be a tedious business, but Red Hat have tried to minimise this



## Software details

What's in the box?

Kernel 2.4.18	KDE 3.0
Glibc 2.2.5	KOffice 1.1.1
XFree 4.2.0	Mozilla 0.9.9
GNOME 1.4.0.4	StarOffice 5.2



The Red Hat installer is flexible and mostly easy to use. It simplifies package selection by offering installation classes.

with a multi-level system. First, you have to choose an installation class from either Workstation, Server, Laptop or Custom. Packages are then organised into task-oriented groups, such as, for example, X Window System, KDE, GNOME, Development, Web server, DNS server and so on. Remember that the same installation CDs are used for both Red Hat *Personal* and *Professional* – so, even though *Personal* is aimed at workstation use, it includes all the server packages as well. Your choice of installation class will determine which package groups you can then choose to install. So, going with Workstation just offers desktop-related package groups, while opting for a Server install presents server-oriented package groups, but no multimedia ones. A custom install lets you select from the complete set of package groups. Whatever your installation class, you can then optionally refine your selection by selecting or de-selecting individual packages by name.

Overall, this package selection system works quite well, except when you need to create a stripped down installation. When choosing individual packages, the installer will check and maintain dependencies for you automatically. A problem here is that the dependencies tend to be a little too strict. For example, you have to install the *Glide* package when you install X – no matter whether you have a Voodoo graphics card or not. Another complaint I have is that you are given little or no guidance as to what the package groups actually contain.

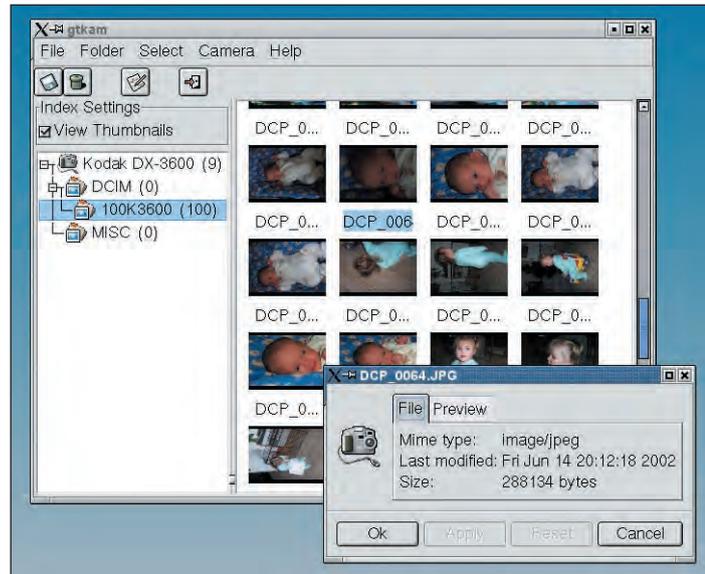
I don't wish to sound negative about the Red Hat installer, because it is certainly isn't bad. It's just that it could easily be better. In fact, it does have many commendable features. For example, Red Hat offers the boot loader *GRUB* by default – a big plus in my book, since I find it infinitely more flexible than *Lilo*. The *GRUB* section in the installer makes it simple to add entries to boot any other operating systems you may have installed and also supports setting up *GRUB* passwords. Similarly, configuration of the built-in firewall is simple and straightforward. You can choose a default security level of high, medium or none – or opt for a custom set-up. This last lets you choose which services are accessible through the firewall, or flag an interface as trusted to let network traffic through unimpeded.

### A choice of desktops

Red Hat now offers a choice between KDE 3.0 and GNOME as desktop environments. While it's great to see that Red Hat have got over their antipathy towards Qt, there are a few problems with their build of KDE. First off, KDE 3.0 was actually a rather immature release (hence KDE's rather speedy release of 3.0.1), so stability is not as good as it should be. Secondly, Red Hat provide a vanilla build of KDE 3.0: they've made no attempt to integrate other software into the KDE environment. Thus you get the standard KDE menus only, and, while GNOME software is available via a sub-menu, tracking down any other software on your system can be a real



The Control Panel provides a central point for launching all the system configuration tools.



Red Hat 7.3 includes multimedia software like *Xine* and *GPhoto2*.

## Alternative software?

New tools to manage software choices

One innovation in Red Hat 7.3 is the adoption of Debian's *alternatives* system to allow the simultaneous installation of multiple pieces of software which fulfil the same function. This system lets the administrator choose which of these installed 'alternatives' to use by default and provides a well-known interface via a system of symlinks. So for example, in Debian, the command `/usr/bin/pager` is actually a symlink – maintained by the *alternatives* system – to the system's default pager, whether *less*, *more*, or something else.

The user can invoke this via the command `pager`, without worrying which alternative is installed. Red Hat don't make much use of this system yet, but in Red Hat 7.3 it does enable you to install both *Sendmail* and *Postfix*, and access your chosen default MTA through the `/usr/sbin/sendmail` command.

The great advantage here is that to switch MTA, you no longer need to uninstall your old one. Your choice of printserver, whether *CUPS* or *LPRng*, is similarly managed by the *alternatives* system.

pain without any menu entries.

This can be best illustrated with an example. Red Hat provide a daemon for auto-mounting CD-ROMs which runs in the background polling for disk changes and automatically mounts newly inserted discs. This also supports an auto-run function which can automatically start a script from a CD when it is mounted. This feature is used on the Red Hat install media, for instance, to launch an appropriate RPM front-end to install the software on that disc. This is a nice touch for new users, and this method of auto-mounting is much more robust than Mandrake's preferred supermount system. The problem is that many will find the auto-run function annoying and immediately want to turn it off. There's a configuration tool available in the GNOME control centre to do this, but when you are running KDE, it's not

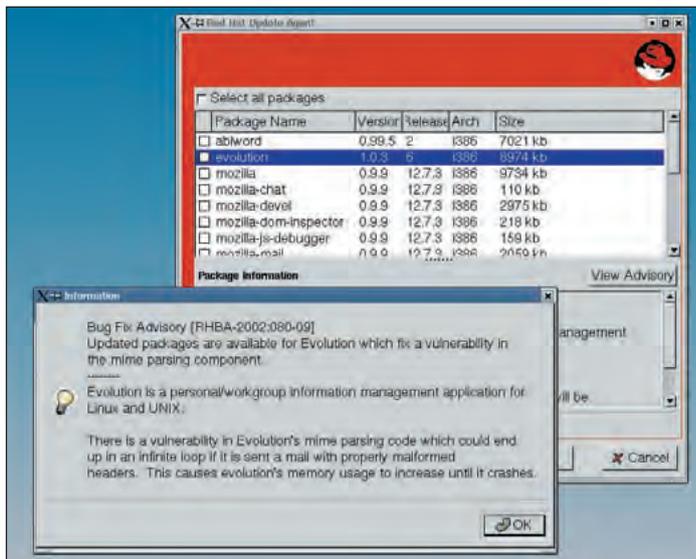
immediately obvious if or how you can do it, since this configuration tool is buried in a cascade of menus.

Red Hat 7.3 contains a wealth of software on its three install discs, so most of your software requirements should be taken care of. This release contains a new focus on multimedia, with packages like *Gnome Meeting*, *Xine*, *GPhoto2* and *dvdrttools* (which provides support for writing DVD-R, DVD+RW, DVD-RW discs) are included for the first time. The only real disappointment on the application front is with the choice of office software. Red Hat have said that they won't bundle the excellent *StarOffice* 6.0 with their distro. More fool them, I say. Thus in 7.3 you have of choice of the under-powered *KOffice* or the clunky, old *StarOffice* 5.2.

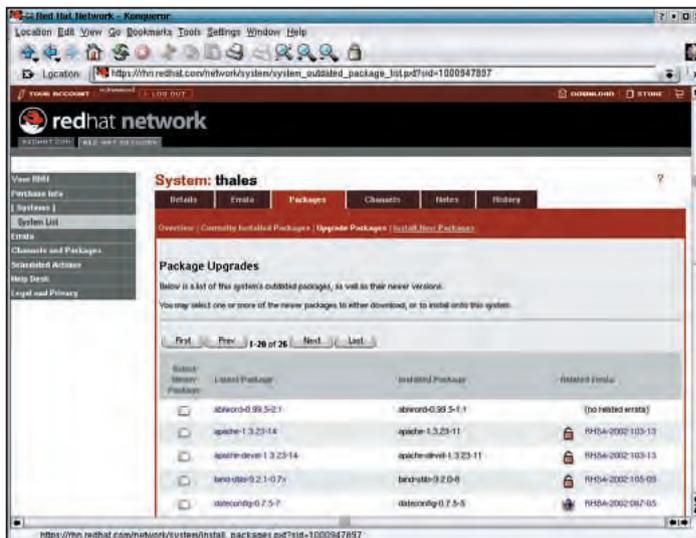
The system configuration tools in Red Hat 7.3 are all available from a



# LinuxFormatReviewsRedHat7.3



The *up2date* tool makes keeping your system up to date with security and bug fixes simple.



The Red Hat Network provides a remote web-based interface for managing the software installed on your system.

central Control Panel; clicking on a tool's icon here will launch it. This provides access to the graphical versions of all the familiar Red Hat tools such as *authconfig* and *serviceconf*. Collecting everything together in a single launch pad like this certainly makes administration simpler, but it's still not always easy to find the tool you want. To configure your network connection do you choose the Internet Configuration Wizard, Network Device Control, or Network Configuration? It's not obvious at first glance. Some pop-up help here to explain what each does is surely needed. What is also still sorely lacking are any tools for doing post install configuration of X. If you want to change your screen mode or depth, or

if you change your monitor or graphics card, then you'll have to resort to manually editing the config file.

## Keeping up to date

Management of software updates in Red Hat 7.3 is performed by the *up2date* tool. This requires that you register with the Red Hat Network before you can use it. Registration is free and gets you a demo level subscription for one machine, which should be quite sufficient for the home user. The overtly paranoid might be put off by the amount of information that Red Hat collects about you and your hardware, however.

*Up2date* checks online for updates to installed software packages and presents you with a list of available

## Red Hat Network

### Deciding on your level of subscription

To use Red Hat's *up2date* tool to manage software updates on your system, you must register with Red Hat Network, Red Hat's web-based support and management system. You don't have to purchase a subscription – simply registering gets you a demo level subscription for one system, which is probably sufficient for most home users. This entitles you to email notification of security and bug-fixes, the downloading of updates via *up2date* (when the servers aren't too busy with paying customers, that is) and remote package management via the RHN website. A basic level subscription

will cost you \$60 per year per machine (a 30-day subscription is included in the *Personal* box), and gives you all of the above, plus priority access to Red Hat's servers for the downloading of updates and ISO images. You also get some extra tools for managing multiple systems, like the ability to view the statuses of all your subscribed systems at once and thus compare them. An Enterprise level subscription brings more tools for managing networks of systems. It lets you group systems by their rôle, manage all systems in a group together, assign administrators, and so on.

bug fixes and security updates, which can then be downloaded and installed automatically. *Up2date* only manages updates: it doesn't install new software. It doesn't offer the kind of flexibility that Debian's *apt* or Mandrake's *update* provide, but it's much more reliable than Mandrake's system. Its does its job of managing updates well, but don't expect it to keep your system at the cutting edge of software releases. Red Hat do only make bug-fixes and updates available, and expect these to peter out as subsequent Red Hat releases are made. One nice feature of *up2date*, though, is that you can actually read the bug fix or security report, before you download a package, so that you can actually decide whether you need it or not.

*Up2date* is only one part of the Red Hat Network support system. Via the RHN site, you can completely manage the software that is installed on a registered system. Here you can view details of your system, including all the software installed, list any updates necessary or even install new software. Updates and new software installs can be scheduled via the site and this works in conjunction with a daemon running on your system. The RHN daemon periodically checks in to see if there are any actions such as updates or downloads scheduled. If so, it will perform them, resolving and fulfilling any dependencies in the process. This system could be prove to be a real productivity boost for busy sysadmins – especially if you opt for the group-based management tools that an Enterprise level subscription brings. It won't be so useful for the home user on a dial-up Internet link.

## Who is it for?

Red Hat 7.3 is a solid, general-purpose Linux distro. It'll do most everything you need, but it just doesn't excite. My problem with this *Personal* edition is that I can't see who it is aimed at.

While it works as a desktop system, it doesn't really have much to offer in this department compared to a lot of the competition. If it's a server system, then what's the point of the Professional edition? The real selling point of Red Hat 7.3 is not the software itself, the supplied manual or anything else provided in the box, it's the good support and management tools that the Red Hat Network brings.

If you just need the software, then rival distros such as Mandrake and SuSE are a much better bet. While a box and a manual give you something nice and physical to put on your shelf, the sort of user that Red Hat will appeal to doesn't need a box – they just need to buy a Red Hat Network subscription. [LXF](#)

## LINUX Format VERDICT

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

Red Hat 7.3 is a reliable, well-rounded, but rather conservative distro. The *up2date* tool and Red Hat Network provide a useful combination for managing software updates, and this is Red Hat 7.3's most compelling feature.

LINUX Format RATING  
 7/10

MS WINDOWS EMULATION SOFTWARE

# Win4Lin 4.0

Hoyt Duff wants it all and NeTraverse delivers.

**Windows emulator straddling the middle ground between Wine and VMWare**■ **DEVELOPERS** NeTraverse■ **PRICE** US\$89.99

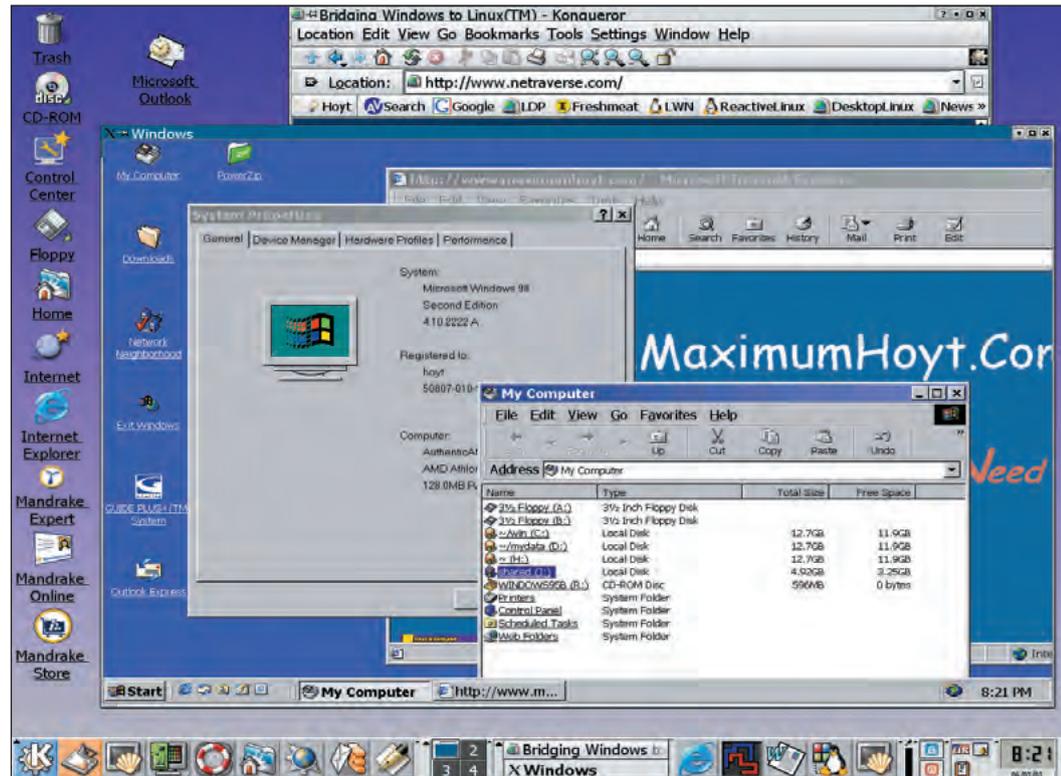
(net download, single user)

■ **WEB** www.netraverse.com

**W**in4Lin is a collection of programs and wizardry that fool Win32 apps into thinking that they are running natively under Windows 95, 95B, 98, 98SE or ME. It fools Windows into thinking that it's the only OS on the computer and into thinking that it is using the FAT filesystem. In reality, all the apps run in an X11 window on your Linux desktop using whatever filesystem you have chosen (*ext2*, *ext3*, *Reiser* and so on) while letting the Linux kernel do the heavy lifting. When used within its limitations, *Win4Lin* is truly an astonishing application.

Nearly two decades ago, *Win4Lin* began life as *SCO Merge*. (The Santa Cruz Operation, you'll recall, was at one time a big player in the Unix game and, curiously, their SCO Unix product began life as Microsoft Xenix.) *Merge* allowed DOS and Windows 3.1 binaries to be run on SCO Unix, making it the grandfather of Windows emulators. That legacy lives on in *Win4Lin* and can be observed in more than a few of its filenames and environmental variable names.

Now, as then, the Unix/Linux hoi paloi ask "Why would you want to do that? Why not just run native binaries?" In the realities of the business world, it is common for some vendors to develop exclusively for the 'wintel' platform as the hardware is inexpensive by commercial Unix standards and the wintel platform is ubiquitous. While many popular consumer applications now have Open Source and Free Software equivalents, most proprietary applications (specialised contact managers, internally-developed programs and so on) still do not offer Linux versions of software.

**It's great for web developers when viewing their work on alternate browsers.**

Let's examine some of the deficiencies of *Win4Lin* that might make it unsuitable for your use. It doesn't handle dongles as well as it should, it doesn't handle USB devices (except mice and keyboards – USB printers are accommodated through *lpd*) and it doesn't work well with applications that require *DirectX*. Forget about *Winsock 2* support or virtual network support for Novell Netware, Token Ring networks, or SOCKS or proxy servers. Don't expect to NFS mount your */home* directory. Don't expect your CD burner to work

or expect you fancy audio manipulation tools to perform. Don't expect to run Windows apps without owning a copy of Windows and don't expect to run multiple instances of *Win4Lin* on a single user license. Even if you do have multiple licenses, don't expect to share any MS Windows code between users as each user must have their own separate Windows installation in their home directory.

## Improvements

On the positive side, you should expect generally faster Windows performance on comparable hardware. You'll also see better stability, the ability to utilise Linux file permissions and direct access to Linux shares from within Windows (if you so choose). You won't have to dual-boot or add a second computer to run those Win32 apps. As a bonus, if *Win4Lin* displays the Blue Screen of Death, kill it and restart Windows very quickly without having to re-boot the computer and halting non-Windows applications.

Two tears ago, installing version 1.0 was nightmare: scripts were broken, the kernel didn't work well and it all had to be installed and massaged by hand. Some images were displayed inverted. Installation is much improved! The GUI installer for version 4.0 is nothing short of amazing as long as you are using a supported Linux distribution. The installer will not only automatically download and install the correct modified kernel and the current libraries, but will check for updates as well and handle the license registration. If you have an unsupported distro or wish to use a custom kernel, there are generic kernel patches available.

Installation is straightforward. Open a root term and run the *Win4Lin-install* binary. It prompts for license information (creating a text file you will later find in */var/win4lin/install*), then checks your running version of Linux to determine if a pre-built kernel is available. If so, it will download and install it. Making changes to your

## Bonus Software

A few unsupported utilities are available from the NeTraverse site: *appwrapper.exe* substitutes for the default Windows shell, allowing a single app to be run without the entire Windows desktop; a HOWTO on using *Win4Lin* with VNC; a Debian kernel packages for version 3.0.

bootloader as appropriate. You must then reboot to load the new kernel. Open a root term and run the installer again. It will check for the most current version of the software, download and install it. (It places the files in `/var/win4lin/installer_tmp` so you can save them and do all this manually if you need to re-install at a later date.) Next, the installer will prompt you to install the Windows CD; if you have a non-standard location for your CDROM, there is a browser to redirect the installer to it. This does not install the Windows OS *per se*, but copies the .cab files to `/var/win4lin` and makes the appropriate modifications to them, then generates boot images for DOS and Windows. Once that is completed, you're finished as root.

As a regular user (it is not advisable to run *Win4Lin* as root) run the installer once again to install the Windows OS for your use. The program creates its virtual "C:" drive in your home directory as `~/win` and also creates a "D:" drive as `~/mydata`. You may manually create up to the traditional limit of 26 virtual drives. Configuration is done both as root and as the user via the same GUI, accessed with the *winsetup* utility.

With version 4, *Win4Lin* addresses some outstanding issues: the addition of Windows ME to the Windows 9.x OSs (it might still support Windows 3.1, but not officially), support for up to 128MB of RAM allocated to *Win4Lin*, support for more swap space, wheel mouse support, better windows resizing, reporting the correct CPU (it used to report everything as a 486), improved software support (including MS *Office XP* and *IE6*, Adobe and

Macromedia products as well as Norton *Antivirus* (it's still MS Windows, remember), as well as improved Asian languages support. Noticeable is the faster startup of Windows 98 in *Win4Lin*. Of great utility is the new support for X C&P, although it may not always work as expected. The kludge of using *xcutsel* is obscurely documented as is the fact that C&P is turned off by default (**Shift-F12/Options/ X Cut & Paste** to enable).

## Future developments

Still awaited by users is improved X Cut & Paste support, improved sound support, better dongle support, improved USB support, as well as *DirectX* and *DirectDraw* support. Will these things materialise? It's hard to tell. NeTraverse doesn't really tout itself to the hobbyist, but instead as a tool for businesses to extend the life of their Win32 apps. While this explains why we may never see *DirectX* support, it doesn't explain why better dongle support and better X Cut & Paste support don't exist since many of those proprietary apps use dongles, and C&P is a boon to productivity.

It will also be curious to see if NeTraverse can engineer any support for the Windows NT/2000/XP series. The underlying design of *Win4Lin* may make that difficult, but it may be a non-issue for a while. We suspect that Win32 apps will have a much longer lifespan than their Win16 predecessors due to slower adoption of 64-bit platforms and OSs.

As for the competition, it is from newcomer Codeweavers with their Wine-based *Crossover Office* product on the low end and VMware on the

## Requirements

### Pre-Compiled Kernels (SMP available)

- Red Hat Linux – 6.0-7.2
- Caldera – 2.3-3.1
- SuSE Linux – 6.4-8.0
- Mandrake Linux – 7.1-8.2
- TurboLinux 6.5
- Connective 7.0

### Hardware Requirements

- Pentium CPU
- 32MB RAM
- 155MB disk space
- OSS-compatible soundcard

high end. All three products have distinct features and significant price differences. All three are worthy of consideration because they do serve different, if overlapping, functions.

The included documentation is beginning to show its age. Containing references to deprecated features, the docs focus excessively on command line configuration (although this does offer a finer grained control of *Win4Lin*'s more esoteric features), and a general lack of focus likely brought on by as series of too many small revisions. Have a look at the Codeweavers documentation as a good example of modern, user-focused information. The *Win4Lin* release notes, however, are well written and extremely useful ([www.netraverse.com/support/docs/400\\_relnotes.php](http://www.netraverse.com/support/docs/400_relnotes.php)) as is the troubleshooting guide.

Version 4.0 also brings an end to the cheap upgrade path along with an increase in the full retail price to US\$89.99 (US\$99.99 boxed). Existing users were allowed to purchase a new license for 4.0 at

US\$49.99 for a limited time (version 1.0 was, I recall, about US\$30). Codeweavers *Crossover Office* weighs in at US\$54.95 and VMware at a hefty US\$329.00. For its intended market, *Win4Lin* provides ample vfm.

For consideration as new features, we would like to see fully functional out-of-the-box cut and paste between Windows and Linux, a script to save the Windows CD installation and re-install it without have to load it again, a script to archive the user's installation and easily restore it as well as another script to archive the root-installed files and their symlinks rather than forcing a complete re-install from scratch, and perhaps a handy desktop icon to launch *Win4Lin* along with some nice icons for our own use. Direct vendor support for the *Win4Lin* kernel hooks would be nice, but the improved cooperation with and from major Linux distros is very welcomed. A *Win4Lin*-enabled kernel available on the same day of the new distro version release is very nice, indeed. As well, the "Backstore" option must be enabled in X11; adding that to the install script would be a nice feature. [LXF](#)

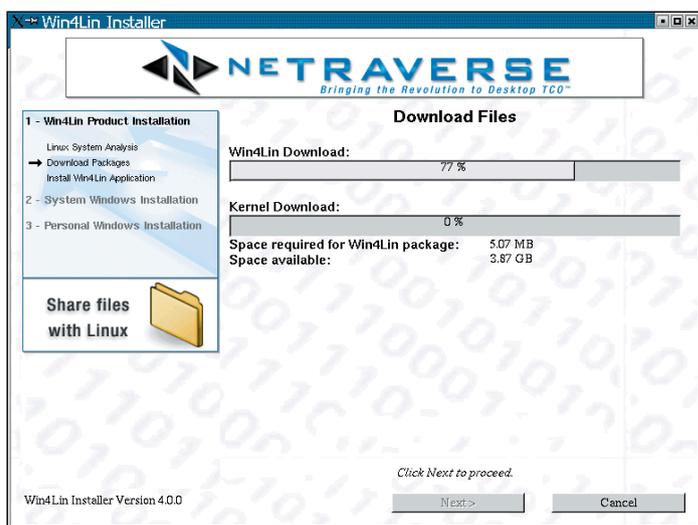
## LINUX Format VERDICT

Installation	9/10
Features	9/10
Performance	10/10
Value for money	10/10

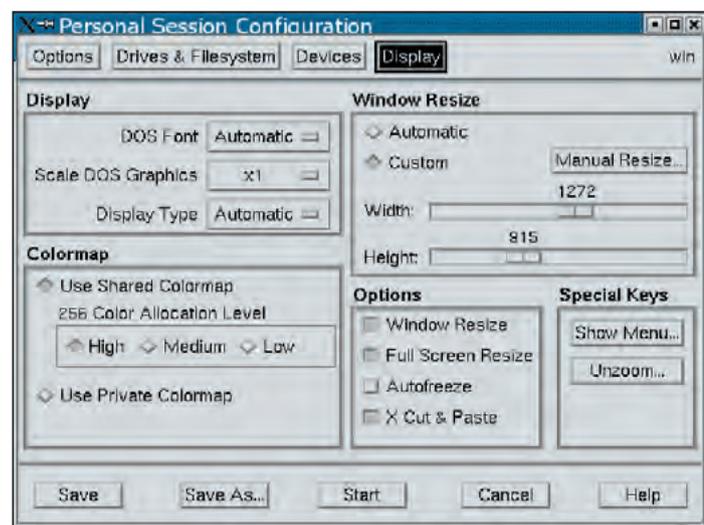
Some deficiencies in specialised areas. Currently the most reliable solution for running Windows 9.x apps.

## LINUX Format RATING

9/10



GUI installer does the heavy work but command line options are available.



Winsetup provided a GUI to configure Win4Lin behaviour.

DISTRO

# Mandrake 8.2 ProSuite Edition

**David Cartwright** unwraps the cellophane on Mandrake's latest monster box and lets us know if it retains its newbie-friendly crown.

**Kitchen sink distro, easy to install. Main rivals SuSE and Red Hat, see also Debian Woody.**

- **DEVELOPERS** Mandrakesoft
- **PRICE** £129
- **WEB** [www.mandrakelinux.com](http://www.mandrakelinux.com)

**M**andrake seems to be generally regarded as one of the easier flavours of Linux to get to grips with – particularly if you read what Nick writes elsewhere in Linux Format. It's not one I've ever played with before, my production systems run Red Hat, though I've installed SlackWare and Debian for clients over the years, so I jumped at the chance to have a play.

## What you get

The ProSuite edition contains DVD containing the workstation installer files, two CDs for the "Server Installer" (for those building a server machine instead of a desktop one), three CDs of commercial applications (paid-for versions of software packages that run on Linux), two CDs of installer sources (source files for the various programs that form a Linux distribution) and finally a docs CD containing 500MB of reference documentation in HTML and PDF, and in a number of languages (English, French, German, Italian and Spanish). We chose the workstation installation DVD for this test.

## Installer

The installer has 20 steps, which are listed down the left-hand side of the screen so you can see how far you are through the overall process. Green items are complete, red items haven't



been started yet, and the orange item is in progress.

The installer seems to have an amount of intelligence – for instance, the PC used for this test has a US Robotics WinModem in, and Mandrake told me that it had found software for which a driver wasn't available but that I could look on [www.linmodems.org](http://www.linmodems.org) for more information. It was also nice to see that my unbranded wheel-mouse was supported – the installer even gave me a little test app to try it out to make sure the driver worked OK.

The disk partitioner is nicer than Red Hat's *Disk Druid*. It's very graphical, shows you clearly what volumes and partitions are there, and makes it much easier than anything else I've seen to avoid accidentally dumping the Windows 2000 and Windows XP partitions that sit on the same disks. It calls ext3 partitions "Journalised", which I guess makes sense to non-experts but it threw me a bit because thanks to the Linuxes I've run in the past, I'm used to seeing 'ext3' instead of something human-readable!

The package chooser has lists that go on forever. The workstation version of 8.2 ProSuite comes on a DVD, which means there's more than

enough stuff on there to fill up the 3.9GB

partition used for this installation.

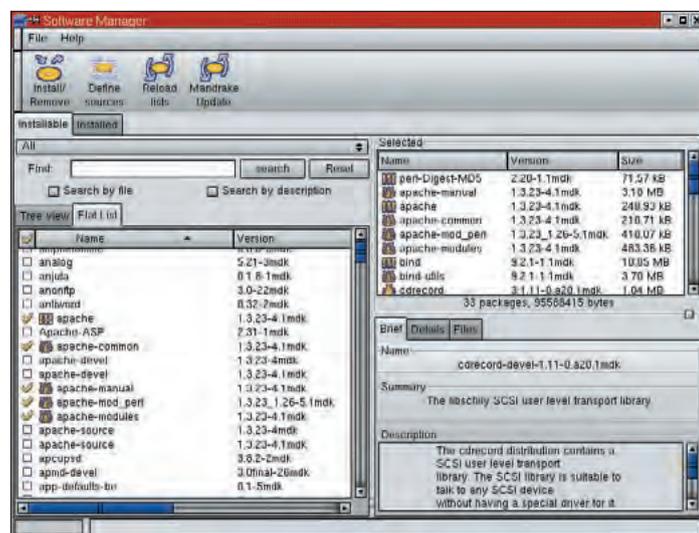
Once you've chosen your packages, go put the kettle on, have a bath, take the dog for a walk, do the shopping, and so on – when it told me "Time remaining: 1:24" I thought "One minute 24 isn't at all bad", until I realised it meant almost an hour and a half!

Once all the files have been installed, you get to choose which to start at boot time. This is handy, as you can choose to install stuff you're not going to use yet, but to not start it until you're ready. When you've told it what to start by default, it lets you choose a

boot loader and tell it whether to include other OSes (in case you have Windows floating around on the same machine) before launching into the X-Windows setup routine. The X installer discovered my unbranded monitor's capabilities with no problems at all – unlike many other Linux installers – and so the X setup process only took a minute or so. Once X is configured, you're given the chance to pull down updates for installed packages from the Web, to make sure the most recent version of everything is installed. Once this is done (and it can take a while – in my case there was 150MB to download) you whip out the CDs and the system reboots.

## Up and running

Once Mandrake has started and you've logged in for the first time, you're invited to personalise your profile – email settings, desktop background



**Software Manager** shows you what applications you have installed, and lets you know which ones have updates available via the Internet.

and so on. This is a quick pick-from-a-menu process, after which you'll see some kind of GUI-based front end (whose nature depends on what you chose in the personalisation process). We chose most of the default options and it all looked pretty bearable.

On the system maintenance side, one of the packages you'll use most is the *Software Manager*. This is the equivalent of Debian's *apt-get* or Red Hat's *update*, and allows you to pull down updates to Mandrake and its associated apps as new versions become available. We found it easier to pick programs from a flat list than to have it presented by app category, but that's just a personal preference. Also under the system maintenance category are the various user interface options (customising your current display manager or picking a new one), network configuration, task scheduling, user management, and the startup manager (which wouldn't run on our test machine – oops!).

Management of the way the system works is done from the *Control Center* [sic], which has options for pretty well anything you can think of – filename associations, hardware and software information, a huge pile of look-and-feel options, networking and peripheral configurations, and even the kernel configuration screens that will look familiar to anyone who's ever done a **make xconfig** in `/usr/src/linux`.

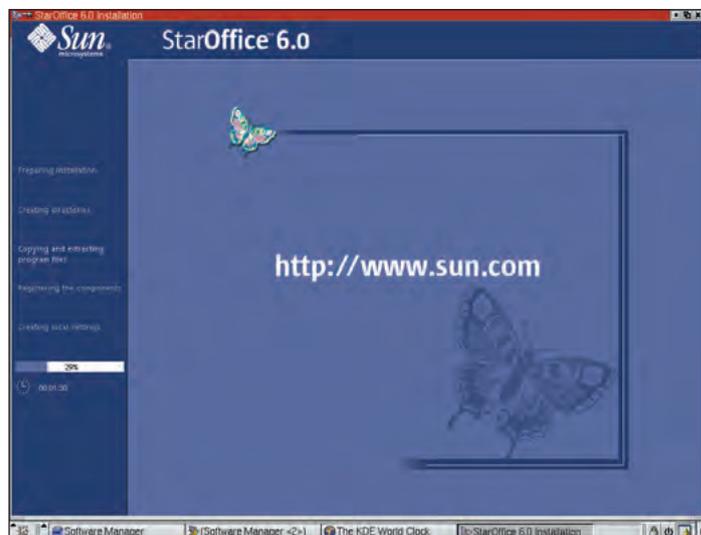
## Application software

System config apps are pretty similar across the board. As far as user programs go, though, Mandrake seems to have gone out of its way to push

Linux down the desktop route – something I've not yet been convinced is a huge market for Linux, but I'm coming round to it. Alongside the pile of games you can install as standard, there are lots of desktop toys included with 8.2. Most prominent is release 6.0 of Sun's *StarOffice* – an extremely usable and versatile office package that is seeing increasing prominence on desktops worldwide. It's a little disconcerting the first time you run *StarOffice* because instead of running the app it fires up the installer, the installation process is merely a one-off 5 min file extraction exercise, however.

Big, multifunction applications aside, there's email (a variety of client software, including *Kmail* and *Evolution*), Web access (the obligatory *Mozilla* and *Lynx*), Usenet News (via *Mozilla*) and IRC. Various IM services are available thanks to the excellent *Everybuddy* multi-service chat application which makes one of my personal can't-live-without features available on a non-Microsoft desktop.

So how did I find my first look at Mandrake? Pretty good, actually. I was very impressed with the installer, which is genuinely easier for non-experts to understand than the Red Hat one (and a whole lot easier than Debian) and in these days of bootable CDRoms you don't even have to go through the palaver of making a boot disk before you can get going. I was pleasantly surprised by the fact that choosing 'Install everything' told me that the 3.9GB partition I'd chosen wasn't going to be enough, though there is a concern in this case that a non-expert might be daunted by the sheer choice



The first time you run *StarOffice*, the installer fires up. Don't worry about getting the CD out, though – it's all on the hard disk already.



When you've picked the packages you want to download, *Software Manager* goes away, fetches them, and installs them for you.

## KEY SOFTWARE COMPONENTS

<b>Kernel 2.4.18</b>	<b>KDE 2.2.2</b>	<b>MySQL 3.23.47</b>
<b>Glibc 2.2.90</b>	<b>Mozilla 0.9.8</b>	<b>Perl 5.6.01</b>
<b>XFree86 4.2.0 (with 3D hardware acceleration if your graphics card can support it)</b>	<b>Apache 1.3.23</b>	<b>PHP 4.1.2</b>
	<b>BIND 9.2.1</b>	<b>PostgreSQL 7.2.12</b>
	<b>Postfix 20010228</b>	<b>RPM 4.0.3</b>
	<b>GCC 2.96</b>	<b>Samba 2.2.3</b>

of stuff available. There's plenty of help and documentation, both installed as standard and available on the separate documentation CD, and the software manager's download facility just works out of the box (the biggest annoyance with Red Hat is having to register with the update service and set your entitlements correctly before trying to update). Finally, it's good that you can choose more than the "approved" set of packages if you wish; where many update systems simply give you the latest recommended version, I could have chosen, through *Software Manager*, to shift up to GCC 3.x instead of sticking to the better-known 2.x.

## Teeny issues

The only problems I found with Mandrake were minor. Some of the application versions aren't exactly up-to-date (the *Postfix* mail server package, for instance, is a few versions old and there isn't an update available via *Software Manager*), and one or two simply didn't work properly (such as the startup program manager, which just threw an error and wouldn't fire up at all). Additionally, it wasn't immediately obvious why the system simply hung at the first login until I

scrolled down the little Web browser screen that was asking me to register my profile and saw there was a button I had to click in order to continue. It must be stressed, though, that these are all minor points – in general the test was trouble-free and it was always easy to find the facilities I was looking for.

In short, Mandrake 8.2 ProSuite is an excellent Linux distribution that lives up to its reputation of comprehensiveness and ease of use, and which will be attractive both to non-expert users and those with complex requirements. **LXF**

## LINUX Format VERDICT

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

Mandrake is a usable Linux distribution with a vast amount of features and applications as standard, which is clearly leading the drive to make Linux acceptable and usable on the desktop.

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

## MINI MAINBOARD

# EPIA mini-ITX motherboard

Nick Veitch checks out another piece of hardware claiming to revolutionise computing.

This mini-ITX motherboard is pretty much unique!

**MANUFACTURER**

VIA Technologies

**WEB** [www.via.com.tw/en/](http://www.via.com.tw/en/)

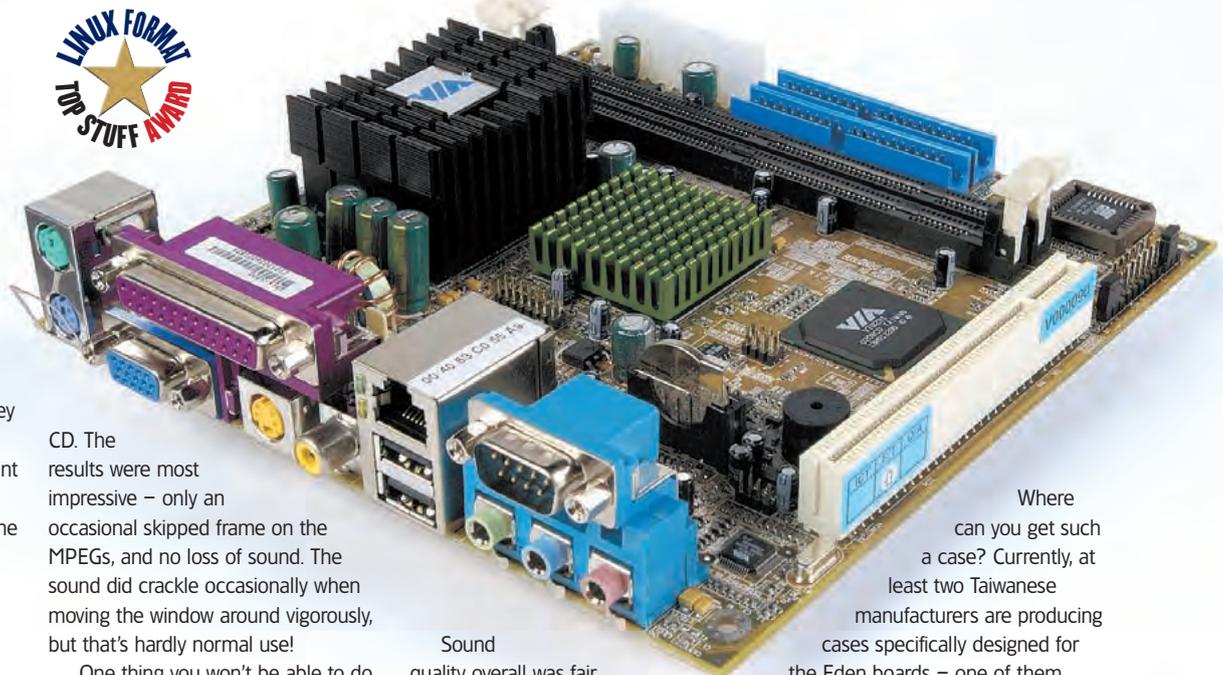
[Vinternet/mini\\_itx.jsp](http://Vinternet/mini_itx.jsp)

**DISTRIBUTOR** [rapid-ltd.co.uk](http://rapid-ltd.co.uk)**PRICE** £ TBC (under £100)

The EPIA motherboard is the first of VIA's integrated Eden boards, which, as you won't be surprised to hear, they have promised will revolutionise computing. But, they may have a point – this tiny board hardly draws any power, fits in a shoebox, and packs the processing power to run all of your standard desktop applications.

The onboard graphics, an embedded Trident Cyberblade isn't exactly state of the art, but is good enough. Considering the chips share the main memory for the graphics, it actually keeps up quite well. Usually such schemes show themselves up in Linux when you start moving windows around, or opening lots of windows on the same screen, but the Via coped just as well as our other desktops.

We tested it playing back MPEGs and DIVX files from the hard drive, while simultaneously copying files to various parts of the same drive from



CD. The results were most impressive – only an occasional skipped frame on the MPEGs, and no loss of sound. The sound did crackle occasionally when moving the window around vigorously, but that's hardly normal use!

One thing you won't be able to do with this board is play 3D games, or use any 3D software for that matter. The processor/graphics combination probably isn't up to playing more recent 3D games, but the lack of any 3D implementation for Linux on this chipset rather scotches any plans for more restrained use of 3D too. The XFree drivers do support a video overlay method for this chipset, so playing movies, etc should not be too much of a problem.

Sound quality overall was fair, which is about the best you get from an onboard chipset anyway. The via82cxxx\_audio modules provides the drivers for this chipset, though on our test installs, Red Hat 7.3 seemed to have some difficulty setting up the various sound services, but we also tested the board with Mandrake 8.2, on which everything worked fine.

One of the interesting onboard features is TV Out. Like quite a few chipsets, this works in an either/or mode with the main display. A connection is detected at boot and the display switched to TV Out if present. Currently the VIA tools for TV Out are not available for Linux, but the TV Out mode may still work if you can configure X in a compatible mode, but after much trying, we gave up on the experiment. VIA did release a driver for X that does support this, but it is not currently available.

## A good case

The EPIA Mini-ITX is, unsurprisingly, designed for the mini-ITX form factor – a standard created by Via themselves. This is smaller than microATX, – the board measures 170mm square, and opens up a range of possibilities for usage.

Where can you get such a case? Currently, at least two Taiwanese manufacturers are producing cases specifically designed for the Eden boards – one of them even runs from a 12v power supply.

If you can't get hold of such a case, the EPIA will fit in most standard ATX cases, though this might be missing the point a bit. One of the bonuses of fanless operation is that everything is quiet – unlike a300W fan assisted PSU.

This motherboard makes a useful base for various computing applications. It's performance is good enough to form the basis of a cheap desktop system, but features like silent operation and low power use, coupled with its tiny size make it attractive for a number of embedded applications. VIA should update the Linux drivers, but with the exception of the TV Out, all the functionality of the board can be realised under Linux. [LXF](#)

## Specs

VIA EDEN ESP 5000 processor	1x parallel port
VIA Apollo PLE133 North Bridge	1x 9-pin serial
VIA82C100A Ethernet (Rhine 10/100)	Line in/Line Out/Mic jacks
Integrated 2XAGP Trident Cyberblade	1x PCI slot
2 x USB	2x ATA 100/66/33 headers
TV-Out	2xDRAM slots (PC133/100)
PS2 Keyboard/Mouse connectors	Award BIOS

## Eden Processor

The processor included with this motherboard is the Via Samuel 2 "Eden" processor, clocked at 533MHz.

Perceived processor power is reasonable for a processor at this clock rate. The test system registered 1064 bogomips. The main selling point of this processor is exceptionally low power consumption (5W) at a 1.2V core voltage. As well as making it suitable for a variety of embedded applications, it means the processor runs without a fan (in fact, the integrated heatsink rarely gets warm)!

## LINUX Format VERDICT

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

A feature-packed, great value board with exciting possibilities.

LINUX Format **RATING**  
10/10

## WEB FILTERING

# SuperScout Web Filter VS

Neil Bothwick examines an enterprise level web filtering solution from SurfControl.

**SuperScout Web Filter VS – Enterprise level web proxy, filter and bandwidth control.**

■ **PRICE** £995 for 50 user licence

■ **PUBLISHER** SurfControl plc

■ **WEB** www.surfcontrol.com

Companies can lose a lot of money from employees using the Internet for anything other than work. This cost falls into two main areas. Bandwidth intended for company business can be tied up by large downloads or streaming media (which may or may not be a problem, depending on the amount of bandwidth available). The other cost is lost productivity, employees participating in mailing list or IRC discussions, or browsing the web, when they should be working. This may not always be intentional avoidance of work, we all know how easy it is to get sidetracked when using the web. A couple of clicks and you've just spent half an hour reading about fascinating things you weren't even looking for.

*SuperScout VS*, from SurfControl, is a way for companies to control their employees' surfing. This isn't the place to discuss the moral implications of this type of control, only to look at one possible way of addressing the situation. *SuperScout VS* is more than a simple URL filter, it provides several ways to manage your bandwidth, from prioritising sites to filtering out banner ads and popup windows. It is designed to fit into any existing network configuration. The most common setup would be to place it between the network and the proxy server, although it is capable of caching web content itself if you don't have a proxy.

## Four million URLs

*SuperScout VS* comes with a database of four million URLs split into categories, so you can filter out large chunks of the web simply by ticking the boxes for categories you want to block. These categories can also be prioritised. The higher priority

```

neil@marvin.digimed.co.uk: /home/neil - Shell No. 2 - Konsole
Session Edit View Settings Help
Note: To use the graphical interface you will need a X-Window server and JRE version 1.2.2 or higher. If you do not have this JRE version installed, please abort this installation and install JRE first! You can continue the installation afterwards by invoking the Configure script.
If you don't want to use the GUI skip this installation step by typing "none"
/usr/lib/jre-1.4.0/bin/java
Creating start script for SuperScout Administrator: admin
Creating start script for SuperScout Remote Watch Tool: watch

If you want to use Bandwidth Control please specify the following settings:
Maximum amount of concurrent network connections (default 200)
20
Bandwidth shaping factor (modulates the impact of the priority setting)
Choose a higher value for more impact (default 50)
100
Setting up registry ...
Successfully wrote 242 keys
Setting up automatic startup/shutdown scripts
Setting up demo license ...

SurfControl Register Tool
Copyright (c) 2001 SurfControl plc.

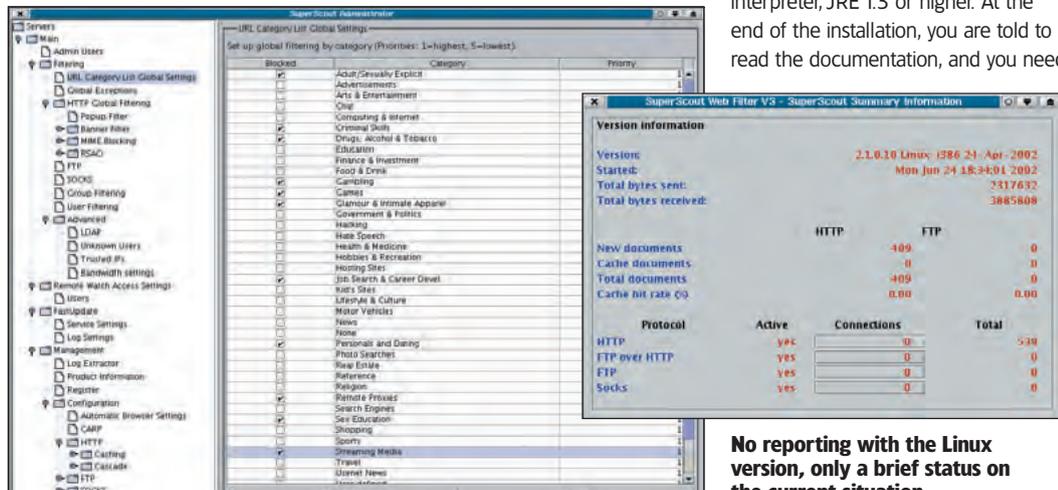
Product registered successfully. Product number: 7

Initialization finished.

Installation complete.
[root@marvin neil]#

```

Installation was simple and quick, and config problems were minor.



The range of configuration options is impressive. Here categories are selected for blocking.

categories get the lion's share of the available bandwidth. This is a way to prevent essential tasks, such as viewing news sites, from being disrupted by large downloads from less important areas. It can also filter according to MIME types, for example to prevent the viewing of bandwidth (and time) wasting content like streaming video.

RSACI is a voluntary system for webmasters to rate the level of adult content on their own pages. *SuperScout VS* can filter these out at whatever level you set. Of course, it's not just about stopping people

displaying adult content in the workplace, you can control other sorts of unsuitable use, such as access to recruitment sites.

At each stage you can specify "URL Exceptions". These are sites or pages that are exempt from the rules. You may find that a site your staff need to access falls within one of the categories you wish to block. The Global exceptions page allows you to assign a different category to this site, or you could just set it to "allow". Equally, you can block individual sites here, even if they aren't covered by the general URL database. As the web

changes so rapidly, a static database would soon become obsolete. The database on SurfControl's server is updated daily and *SuperScout VS* is able to download and apply the updates automatically.

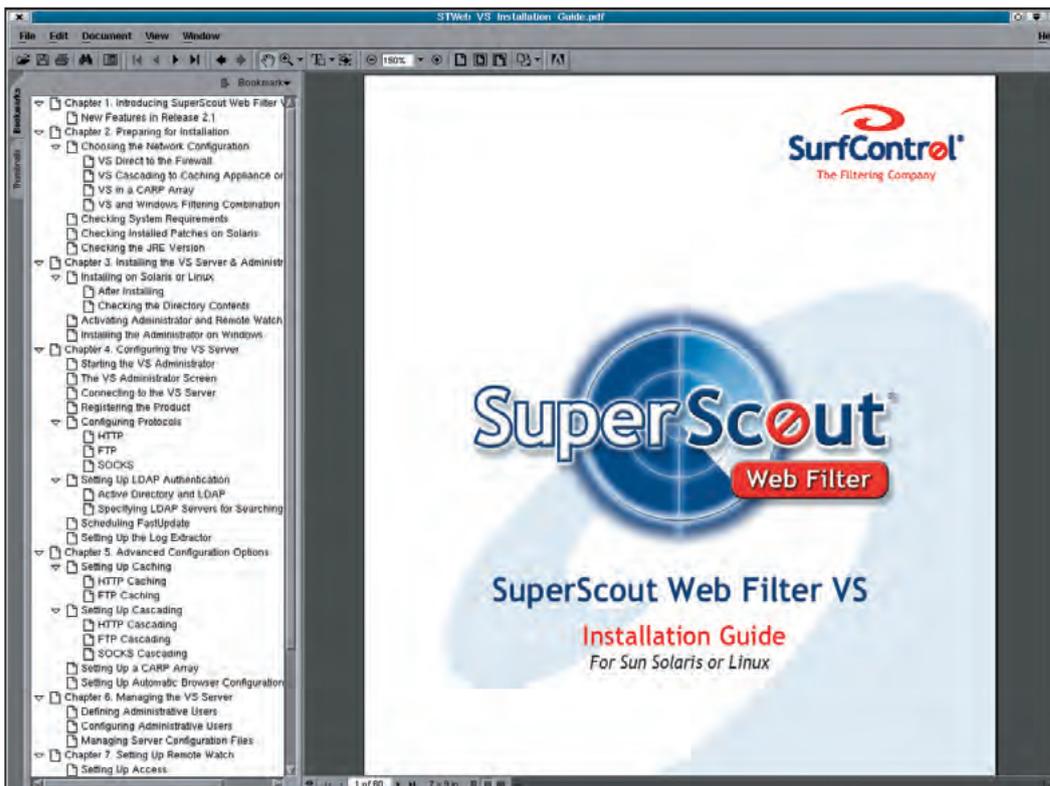
## Installation

The minimum system requirements are a 500MHz processor, 256 MB of RAM and a 2.4 kernel. Installation was simply a matter of running a shell script. The first time I ran this script it failed. It needs *uncompress*, which wasn't present on my Mandrake 8.2 system. A quick call to *urpmi* and it was time to try again (the package you need is *ncompress*). Apart from that glitch it installed correctly and easily. *SuperScout VS* also requires a Java interpreter, JRE 1.3 or higher. At the end of the installation, you are told to read the documentation, and you need

No reporting with the Linux version, only a brief status on the current situation.

to. There are a lot of options when configuring *SuperScout VS*, the install guide is a 60 page PDF file, with another 113 pages in the user guide.

The initial configuration was straightforward. After setting up an admin user and password you can set up caching or proxy servers. There are several options for either, including specifying domains or URLs that are not to be cached. It is also possible to specify multiple proxy servers to share the load. For example, all .com requests could go to one server with another handling the rest. If you use *SuperScout VS*'s own caches, you can specify how long docs are cached for,



The documentation was comprehensive and easy to follow, with few omissions.

and how often the cache is cleaned out. You can do this frequently to minimise the load when it's done, or only do it at night if the network is mainly used during working hours.

*SuperScout VS* can also cache FTP requests, with similar options. The FTP cache can have separate expiry times for different types of file, based on the filename extension. Each browser needs to be told to use *SuperScout VS* as its proxy server. The most transparent way to handle this is to set it up through your DHCP server. It's also possible to use a proxy auto-configuration (.pac) file for automatic proxy setting in the browsers, although *SuperScout VS* doesn't generate this file for you.

## Filtering options

Filtering can be implemented globally, by IP address, by users or by groups,

or any combination of these. User and group filtering requires that users login when first accessing the web in a session, the login information is held in an LDAP database. This is one area where the generally comprehensive documentation falls down badly. It gives no information on setting up a suitable LDAP facility, not even a link to some online resources on this. Admittedly, this is an enterprise product that is likely to be installed by a system administrator familiar with such things, but it does jar with the helpful nature of the rest of the documentation.

IP address filtering should be used with care when addresses are allocated by DHCP, and would only filter by physical machine rather than the user of that machine. One use for it would be where you simply want the global settings to apply to everyone, so

you would set up a "trusted" IP address of the whole network and apply the global filtering to that address. This would avoid people having to log in, but would prevent useful reporting by individuals.

## Linux deficit

So far so good, we have a powerful and flexible way of monitoring and logging web access. However, it falls down when trying to access the reports, at least from a Linux perspective. The program supplied to generate and display reports only runs on Windows NT or 2000. This program runs as a server, *SuperScout VS* can run a log extractor that sends the information to the *Reporter* server. This information is stored in an SQL database and the *Reporter* can then generate reports, either live when connected to with a web browser, or scheduled reports that are saved to disk or emailed.

It doesn't make sense to only make this part of the package work with Windows NT when the rest also runs on Linux or Solaris. You can view the reports on a Linux box, but a Windows NT machine is compulsory to generate them. Linux is more than capable of handling this task. Even simply writing the information to a MySQL or PostgreSQL database

and providing a couple of Perl or PHP scripts to query and display it would have increased the usefulness of this product enormously. There is a *Watcher* program that runs on Linux, but this only monitors the current state of the system, it doesn't provide the sort of information the Windows version of this package gives.

## Comprehensive-ish

*SuperScout VS* is, on the whole, well documented. It is a complex product but the two PDF manuals cover most topics in sufficient depth. It also comes with some ancillary information, such as a booklet on developing a company's Internet Acceptable Use Policy. Although this is aimed as justifying investment in *SuperScout VS*, it does provide useful information, advice and resource lists.

There were a couple of problems when setting up *SuperScout VS*. The first, minor, point was that many paths in the configuration, such as the cache directories, point to /tmp. Not only is using /tmp a bad idea, but the manual specifically states that the HTTP and FTP caches should use different directories. More serious was the failure when trying to configure global URL exceptions. Emails and telephone calls to SurfControl's technical support, who were prompt to answer in both cases, failed to get an immediate solution. Questions on the Linux versions were being forwarded elsewhere. From talking to the company, it appears the move to Linux is fairly recent, and it shows. While *SuperScout VS* is a well thought out and well documented package, the Linux support is lacking. Until the Linux version is as complete as the Windows package, it is hard to recommend *SuperScout VS*. When that day comes, this should be a powerful and useful package. [LXF](#)

## The user's perspective

### Business as usual

*SuperScout VS* is almost totally transparent to the user. He needs to log in if user or group filtering is used, but after that there's no noticeable performance hit. In fact, if you weren't already running a proxy server, it will work faster because of the built in

caching. The only time a user would be aware of *SuperScout* in normal use would be when they try to access a blocked page. Reconfiguration is so simple that the sysadmin could simply remove a block on such a site if the user had a genuine need to visit it.

## LINUX Format VERDICT

Ease of use	8/10
Features	4/10
Performance	8/10
Value for money	5/10

Until the reporting system is ported to Linux, it is difficult to recommend *SuperScout VS* for a Linux network.

LINUX Format RATING  
**6/10**

## WORKSTATION DISTRO

# Conectiva Linux 8.0

Conectiva may be the most popular Linux distro in Latin America, but does it have anything to offer this side of the pond? **Richard Drummond** finds out.

Rival distros include Mandrake Linux, Standard Edition, Red Hat Personal and SuSE Personal.

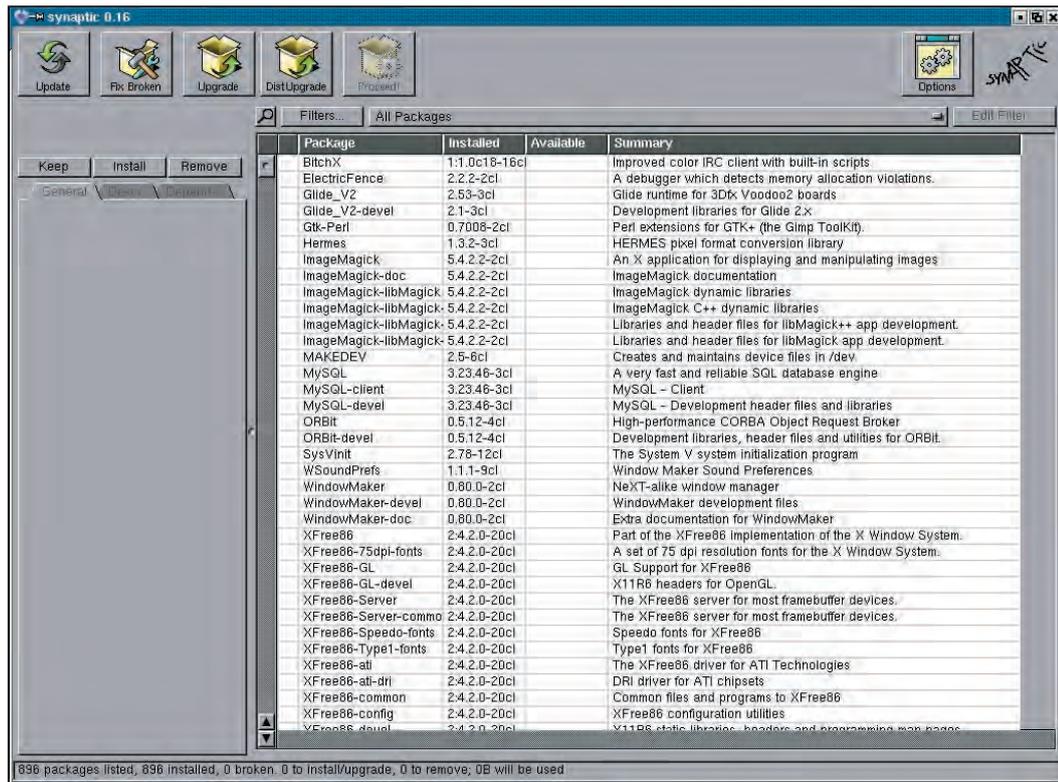
- **DEVELOPERS** Conectiva
- **PRICE** R\$109 boxed (£26)
- **WEB** [www.conectiva.com.br](http://www.conectiva.com.br)

Conectiva Linux hails from Brazil, a country renowned for its world-cup winning football team but equally well known in Linux circles for producing talented kernel developers such as Rik van Riel and Marcelo Tosatti (both of whom are employed by Conectiva). Conectiva Linux is remarkable in that it is the only mainstream RPM-based distro that uses a port of Debian's *apt* (*Advanced Package Tool*) for package management. We were impressed, when we last looked at Conectiva, at how well *apt* eases the management of software updates even on an RPM-based system.

Conectiva Linux 8.0 is available in a boxed edition with two install CDs, manuals and *StarOffice* 5.2. We don't know of a European distributor, however, so we are once again testing the GPL'd download edition. The Conectiva website provides little information on Conectiva 8.0 in English, and, while Conectiva has been translated into Portuguese, Spanish and English throughout, there is a disadvantage of having no Conectiva-specific documentation available in English. Of course, the various software packages that make up Conectiva all have English documentation, and the installer offers online help in all three languages and is actually very easy to use.

## Installing Conectiva

Conectiva derives a lot from Red Hat, and in installation is very similar, albeit with a slicker and more attractive user interface. Hardware detection has improved since release 7.0, although we did experience system lock-ups during graphic card configuration. Like Red Hat, Conectiva isn't happy on multi-headed machines.



You don't have to wrestle with *deselect* or *apt-get* to keep your Conectiva installation up to date or to install new software. *Synaptic* provides a user-friendly face.

## Features

**Kernel 2.4.18**  
**Glibc 2.2.4**  
**Xfree86 4.2.0**  
**KDE 3.0rc3**  
**Gnome 1.4.0.4**  
**Windowmaker 0.80.0**

Partitioning is handled by *Parted*, and is mostly predictable. One fault, though, is that Conectiva performs automatic partitioning by default. You must tick a check box to force manual partitioning, and this is all too easy to miss. (It won't let you back-track and change your mind, either.) On the plus side, ext2, ext3 and Reiser filesystems are supported as well as software RAID and LVM partitions.

Package selection works much like Red Hat with a similar multi-level selection system. It's rather more manageable, however, because during

the install phase you can only install software from the first install disk. Any other software must be added from the installed system.

## The apt advantage

Conectiva offers a broad variety of desktop and server software. Users are given a choice of GNOME, KDE or Windowmaker desktops, and you can install Apache, Samba, OpenLDAP, PostgreSQL and all the usual range of services. Configuration hasn't improved significantly over the previous release, with just the usual range of Red Hat admin tools plus the aging *linuxconf*.

The version of KDE 3.0 supplied with Conectiva is a pre-release, and is noticeably buggy. A nice touch, though, is that Conectiva supply their *Crystal* theme— which feels like a breath of fresh air after so long spent looking at the default KDE theme.

Where Conectiva scores over Red Hat is with its update system. Version

3.0.1 of KDE is now on Conectiva's mirrors and fetching and installing it is a breeze with *apt*. No subscriptions or registrations are required.

It would be easy to dismiss Conectiva as merely a Brazilian version of Red Hat with a prettier veneer, but *apt* – in conjunction with Conectiva's more liberal updates policy, mean that many will find it preferable despite being a less well-known brand. **LXF**

## LINUX Format VERDICT

Ease of use	7/10
Performance	6/10
Features	7/10
Value for money	9/10

Expect minor installation problems, but the combination of RPM's ubiquity with *apt*'s power is an attractive one.

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

# Using SANs and NAS

**David Coulson** finds out how NAS and SAN installations work.

- **AUTHOR** W. Curtis Preston
- **PUBLISHER** O'Reilly
- **ISBN** 0596001533
- **PRICE** £20.95

**W**ith disk space requirements of e-commerce facilities, big business and data centres increasing quicker than many people can possibly fathom, the need for reliable and redundant storage systems is a priority among the administrators, along with the management who have to justify the cost of high-end data storage methods. With the choices being between an IDE hard drive, and a RAID array employing fibre-channel technology, it quickly becomes far more complicated than expected.

In true O'Reilly style, Preston has managed to cover almost every aspect of network-based storage within two-hundred pages, without cutting back

on detail or avoiding more complex configurations which people may need. This book covers both the hardware and software needed to construct a storage area network, rather than assuming knowledge of either or both, so if you're thinking of building a SAN or NAS system, you can start from scratch on the right foot, rather than trying to boot your knowledge in a specific area in order to keep up with the book. Fail-over redundancy, regular back-up, and the appropriate network architecture for use with a SAN or NAS are dealt with, and can be adapted for less critical situations with some ease.

Along with the NAS or SAN itself, the various protocols which can be used to connect to a storage unit by a client are covered in some depth, and their various merits, or lack thereof, are discussed and compared in easy-to-read tables. The book contains many simple and logical diagrams of



networks and storage systems – so gone are the days of trying to visualise what the author is trying to describe in the book. Many of the diagrams and, indeed, the entire book, do expect some level of technical knowledge, and while many topics are covered from a beginner's standpoint, unless you are seriously considering implementing such a system, it is quite likely to be way over the reader's head.

While not being a Linux-specific book, the numerous platforms used for SAN and NAS configurations are looked at, and a variety of vendors which offer kit running the operating

systems have their names dropped.

This is a perfect book for technical people wanting to increase their knowledge before having to deal with SANs or NAS installations, but also works well for beginners.

## Linux Format VERDICT

The definitive guide for administrators and builders of network storage systems and management solutions.

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

# Programming PHP

**Nick Veitch** reviews a book which promises much but doesn't quite deliver.

- **AUTHOR** Rasmus Lerdorf and Kevin Tatroe
- **PUBLISHER** O'Reilly
- **ISBN** 1-56592-610-2
- **PRICE** £28.50

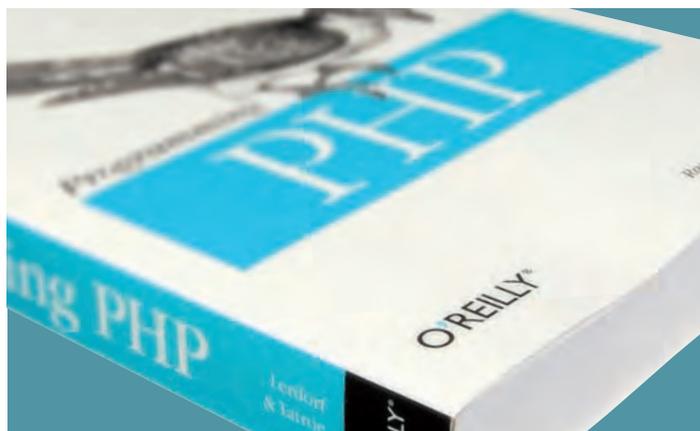
**A**s this book is co-written by the creator of PHP, you'd expect an authoritative guide, and, for the most part, you'll not be disappointed.

The book starts with the basics of the language – how variables are handled, expressions and operators, functions, arrays and objects. Later chapters cover server side scripting basics, using databases and even using the *GD* extension to generate graphics on the fly. There's a section on PDF, and the *lingua franca du jour*, XML.

Within each subject area, concepts are introduced and explained in a

logical order, with plenty of examples for reference. However, some of the topics under discussion are dealt with a little too briefly, and you may be left wanting more. The databases chapter is a good example of this. Pretty much everyone using PHP seriously on the web is using it in conjunction with a data source, and I would imagine the PHP/Linux/MySQL combo crops up more often than not. Surprisingly then, the databases section lasts for only 13 pages out of 470-odd, and even then only covers the *PEAR DB* interface. The examples are fair enough, but only touch the surface of database interactivity (e.g. there's no example shown to check whether the SQL query actually contains any data).

A later chapter deals with securing your PHP code, and although this is only a short section, it includes some wise advice and a few examples of the



potential pitfalls of failing to anticipate the misuse of your code.

Unfortunately, this book seems less than complete. Missing out on the compile-time extensions included as standard (e.g. *MySQL*) means you are going to have to look elsewhere for good info. The function reference for example, only covers the basic PHP functions – not even those that can be specified when building (*MySQL*, for example). So, it isn't going to replace looking up the PHP homepages for documentation.

While this is still one of the better PHP titles available, it could and should have covered all the subject areas a little more thoroughly.

## Linux Format VERDICT

Probably one of the better PHP books available, but seems a bit rushed in places.

LinuxFormat **RATING**  
 **7/10**

# Programming with Qt (2nd Edition)

Maurice Kelly reviews the ideal Qt introduction.

- **AUTHOR** Matthias Kalle Dalheimer
- **PUBLISHER** O'Reilly
- **ISBN** 0596000642
- **PRICE** £28.50

**G**raphical User Interfaces changed the face of computers, making life much easier for computer users everywhere. Except for the people who had to implement them of course, for whom program interfaces became a lot more complicated. It is now a lot easier for programmers with the advent of feature rich GUI toolkits such as Qt, which is the basis for the hugely successful KDE project.

*Programming with Qt* has been updated to cover the new features of Qt3. The book forms a strange hybrid between tutorial and reference manual. The author explicitly states

early on that the Qt reference documentation (supplied with all distributions of the toolkit) forms a more complete reference than any book ever can. This is a fair point, and means that you aren't buying a printed version of freely available docs, but instead are getting more of a general overview, and some useful pointers from a proficient user of the toolkit.

An intro section is swiftly followed by a two chapter tutorial on basic techniques before a "guided tour" of Qt widgets and standard dialogs is thrown in. At this point the readers knowledge of Qt should be sufficient to get stuck into her own apps, using following chapters for more detailed info when the going gets tougher. The author does not forget that fact that Qt is a multi-platform toolkit, referring to Win32-specific gotchas and providing whole chapters on portability and using Qt with *Visual C++*.



There are weighty coding examples throughout the book, and while these are often well-commented, it was disappointing to find that line numbers were not printed, despite the explanations referring to specific lines. The examples were not supplied on a disc with the book, but at <http://examples.oreilly.com/prowqt2/>. A mention must be made of the author's continual attention to good GUI design – as well as a dedicated chapter on the subject, there are plentiful reminders in the remaining text about when it is appropriate to use particular widgets.

This is a very well-written book, and a worthwhile read for all but the most advanced Qt programmers. It requires knowledge of C++, but Qt concepts are brought up from a basic level and are easy to get to grips with.

## Linux Format VERDICT

Considering the freely available reference material this is probably the only Qt book you will ever need to buy.

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

# Web Services Essentials

David Coulson looks at a broad, yet comprehensive, guide to web service languages.

- **AUTHOR** Ethan Cerami
- **PUBLISHER** O'Reilly
- **ISBN** 0596002246
- **PRICE** £20.95

**W**ith many web services being employed to pass data between a server and a client program, rather than just for use with a web browser, a number of markup languages have been developed. Each being slightly different, they all allow a server and a client to pass info back and forth, usually over HTTP, much like any other sort of connection. The most common use of such a language is for a Remote Procedure Call, or RPC, where the client passes a data to a

specific section of a server program, and the client is returned the appropriate results.

This 250 page guide covers XML-RPC, SOAP, UDDI and WDSL; probably the four most popular RPC markup languages on the Internet. Each is covered in some depth, but if you're thinking of using one for development, then you may wish to purchase a book more specific to that language – though much of the info needed to get going, once you have the basic idea behind the language clear in your mind, is on the Net. Many of the worked examples are written using Java, so knowledge of the use of the various mark-up languages with the language being used, such as **XML::RPC** with Perl, will be needed,



as while the actual usage of the system is the same, many parsers and servers behave differently and require the developer to write the code slightly differently.

This book is not supposed to be a thorough guide to RPC using XML, as O'Reilly offer books on each of the languages which have been included in this publication. However, if you're just getting started creating apps which interact with a remote server over HTTP, using any of the languages, then it gets you started with the

fundamentals without confusion or jumping in too deeply without sufficient background knowledge of the procedures used. **LXF**

## Linux Format VERDICT

Useful coverage of many aspects of XML service creation if you're unfamiliar with many of the languages.

LinuxFormat **RATING**  
**8/10**

# HotPicks

The best new open source software on the planet!



**Nick Veitch**

The ed is never too busy to write about the joys of Free Software.



**Richard Drummond**

Our reviews ed maps out a few interesting apps and picks the hottest.

**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

Look at the Stars	35
Firewall Builder	36
PFAEdit	36
LGeneral	37
Parted	37
Generic Mapping Tools	38

## 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!



## PICTURE MANAGEMENT

# Look at the Stars

■ **VERSION** 0.5 ■ **WEB** <http://lats.sourceforge.net>

**L**ook at the stars, is, quite frankly, a completely ridiculous name to give to a very interesting software project. The idea is not a new one – a picture management tool that will show thumbnails of all image files, and permit further operations on them – a media manager if you will.

However, there are few examples of this type of software under Linux, and probably the best known, *Compupic*, isn't open source (and hasn't been updated in aeons).

Using the *GTK* toolkit (and various *GNOME* libraries) the user interface is quite simple and uncluttered. The window is split into three main panes – one for the directory listing, one for thumbnails and one large viewing window for the currently selected image. The view can be changed to fill



**Does this boy have a great future in the Free Software world? Will his first word be Emacs or Vi?**

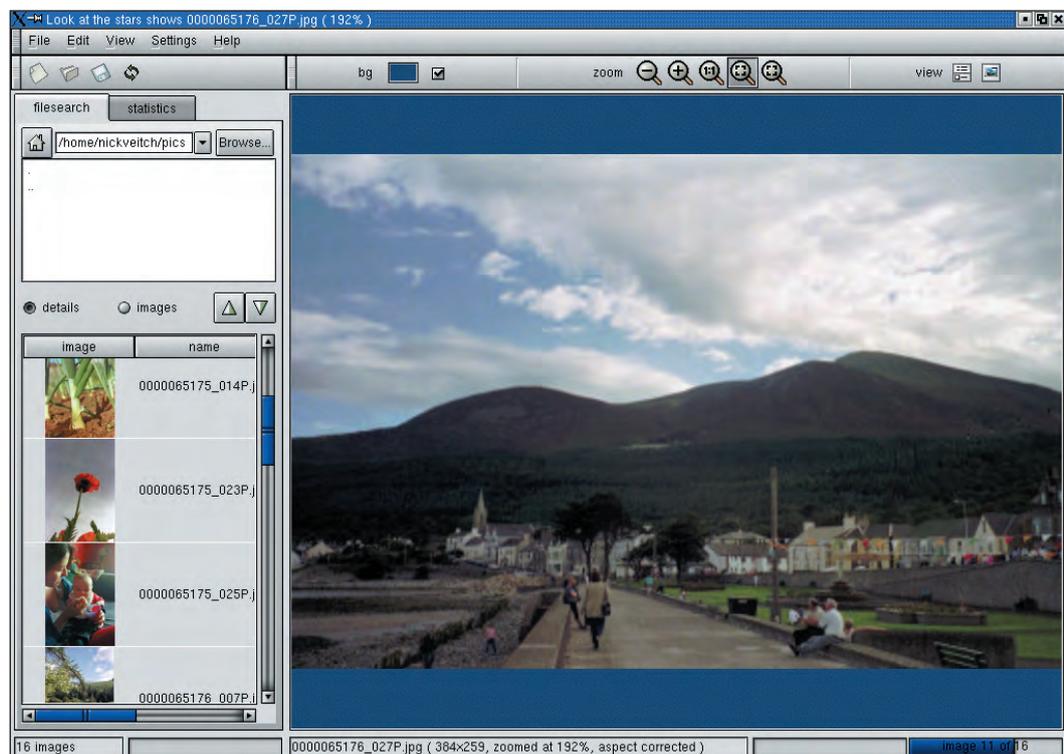
the larger window with thumbnails if you like. Clicking on the image, as you may expect, loads it into the main window, and you can then scroll and zoom in. A few preferences can be set, mainly to do with the rendering of thumbnails and the main image. Thanks to the various *GDK* libraries a large number of file formats are

supported, so your image collection should be viewable.

There isn't much functionality yet – things expected to be added are a slideshow mode and various options for moving and renaming the images.

What's perhaps more remarkable about this project is that it's the work of a single developer, and its his first major project. I'm sure he'd be grateful for any feedback either on the code or on new features and how they might be implemented. As ever there are plenty of ways to help out, so visit the project homepage at sourceforge.

The worst thing about this project is the name! There is nothing in particular wrong with it, but having the binary called *look\_at\_the\_stars* is a bit annoying. I recommend that, like me, you *mv* it to *lats*.



The Look At The Stars (stupid name) interface is simple, but effective and intuitive.

## FIREWALL MANAGEMENT

# Firewall Builder

■ VERSION 1.0.3 ■ WEB [www.fwbuilder.org](http://www.fwbuilder.org)



**Firewall Builder** – behind the daunting GUI lies a great tool for administering firewalls on larger \*nix networks.

**F**irewall Builder is a system to simplify the management of Linux firewalls. It consists of three components: a GTK+-based GUI, which lets you specify your network topology and firewall rules; a compiler, which takes the rules you have defined

in the GUI and compiles them into a script suitable for setting up the firewall; and a firewall daemon, a process which runs on the firewall itself and automates the job of deploying your compiled firewall script. Obviously, this division of

labour means that you don't have to run the *Firewall Builder* GUI on the machine which actually runs the firewall; in fact, it's more than likely that a firewall box will not be running X, anyway.

The Firewall Builder GUI is an object-based system, and the first task is to define the devices in your network or networks as objects in the GUI. While you can do this by hand, *Firewall Builder* has a number of tricks up its sleeve to automate this process. Firstly, it can read the `/etc/hosts` file from your the machine its running on and so discover what hosts the local machine knows about. Secondly, it can import a DNS zone from your network's DNS server, and thus find any network nodes that it knows about. And thirdly, it can you use SNMP (Simple Network Management Protocol) and recursively scan your network SNMP hosts. Whichever method is most appropriate for you, it can certainly speed up the tasks of list of defining your network.

Next you need to tell *Firewall Builder* which of these hosts is the firewalls, describe its interfaces, and define the rules to apply. Again, here,

*Firewall Builder* has a wizard to help out. The build rules wizard can construct an initial set of rules for a firewall to protect a local network, optionally with a DMZ and optionally with masquerading. You just then need to refine your security policy, by dragging and dropping the services you want to let through the firewall and any hosts or sub networks which are to be the source or destinations for these services onto the policy page. When your happy with all your rules, click compile to build and deploy your rules scripts to the firewall. Compilers are supplied to handle *iptables* firewalls under Linux and of firewalls on OpenBSD and *ipf* firewalls on FreeBSD.

*Firewall Builder* speeds up firewall administration, especially for large networks. The GUI is initially quite daunting, but, while the online help is limited, the project home page contains links to some excellent tutorials to get you up and running in no time. It may be overkill for a masquerading dial-up server on your home network, but it's perfect for managing complex networks with lots of services.

## FONT EDITOR

# PFAEdit

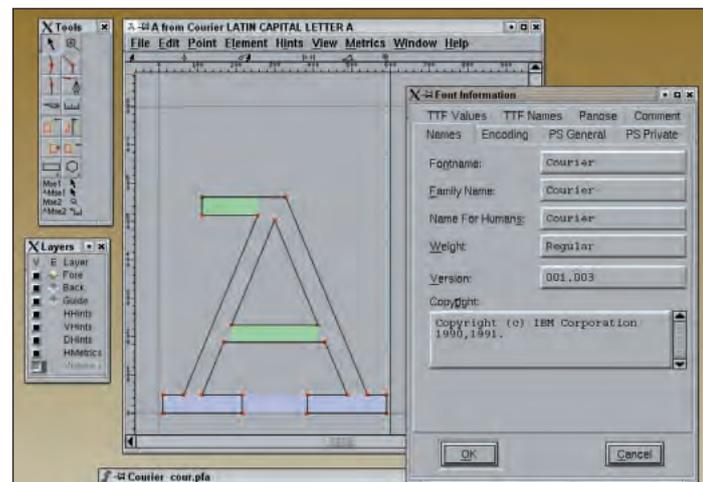


■ VERSION 20020618 ■ WEB <http://pfaedit.sourceforge.net/>

**P**FAEEdit is a sophisticated graphical editor for designing and manipulating PostScript fonts. With *PFAEdit* you can create fonts from scratch or import fonts in a number of formats in any of the various popular encodings such as PFA, PFB or Mac binaries. Although intended for PostScript fonts, it can import and export TrueType and OpenType fonts, too. It can also generate PostScript metric files (AFMs), and rasterise and generate bitmaps at specified resolutions and save them out as as part of a TrueType font or as BDF. The design of composite glyphs such as accented letters and ligatures are supported, and *PFAEdit* can support a wide variety of encodings as well as CID-keyed fonts (a method used for storing PostScript fonts with large numbers of glyphs – such as Chinese, Japanese or Korean characters – efficiently).

The drawing tools that *PFAEdit* supplies for designs glyphs are advanced. These include a palette of drawing tools for creating points, tangents and curves, for cutting splines and for general rotations and scaling – as well as time-saving tools for averaging out or spacing points along a line or parallelising point. You can import a bitmapped-font as a guide for designing your glyphs, or, if you really especially lazy, get *PFAEdit* to trace around them for you with its autotrace feature (although results can vary greatly here). Then there are various wizard-like features. A 'Find problems' function is implemented that will help to spot and correct problems in your glyphs, such as open paths, while a 'Simplify' function exists which will try and remove unnecessary points from you design.

*PFAEdit* has tools for specifying hints in a PostScript font glyph (hints

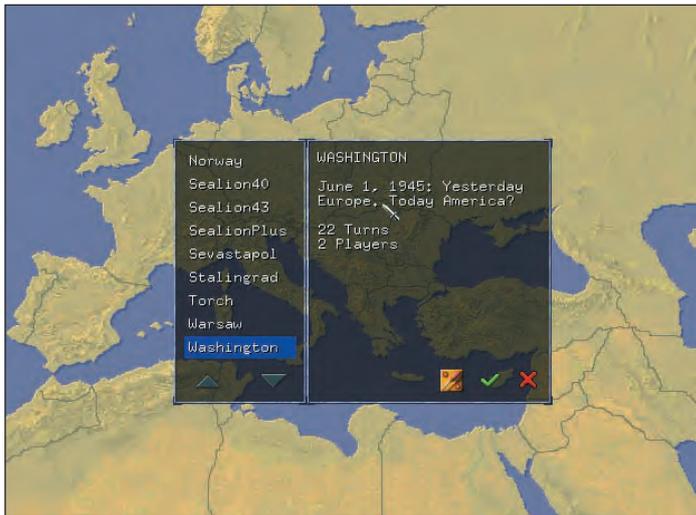


**PFAEdit** is useful tool for editing and manipulating fonts, even if you are not adventurous enough to want to design your own font.

are used to make glyphs readable at low point sizes). It even has an autohint function which will determine and set hints for you automatically. TrueType hints, properly called Instructions, are not supported at the design stage, although *PFAEdit* will try to convert PostScript hints to TrueType instructions.

When it comes to features to aid font design, *PFAEdit* is the equal of many commercial packages. However,

it's an incredibly useful program to have around even if you are not intended to design you own fonts. It can find use as a graphical tool for font conversion – rather than having to mess around with a string of shell commands; or it can be used to modify a font for your own particular use. For example, if you favourite font lacks a certain glyphs, you can use *PFAEdit* to copy and paste them from another font which does have them.



Some of the demo scenarios definitely fall into the 'What if?' category!



If you fail the French can rewrite history and defend the Maginot line.

## STRATEGY GAME

# LGeneral

■ VERSION 1.0.3 ■ WEB <http://lgames.sourceforge.net/>

To those that say that there are no games available for Linux, I say, "Rubbish". You may not be able to pop down your high street boutique to buy the latest titles for Linux, but re-implementations and clones of tons of classics are available. If you rate playability over originality, then it's a great gaming platform

Take *LGeneral* for example. This is a faithful clone of that old DOS favourite from SSI, *Panzer General* – a turn based strategy game set in World War II. It might not be politically correct, and the graphics and sound may be crude by today's standards, but there's gameplay a-plenty to be had.

*LGeneral* is based on the cross-

platform *SDL* library. Source code only is available from the project home page, but it is simple enough to build yourself. The source download contains no games files – scenarios – itself, but a converter is available to translate scenarios from the original game. The demo scenarios from *Panzer General* are pre-converted and available as a download. This features a selection of battles from World War II for you to play, some historical, some fictitious. Several additional scenarios are available for download from the web site, and, of course, there are still plenty of places on the web where you

can download add-ons for the original *Panzer General*. A scenario editor for Linux is also apparently in the works for you to roll your own campaigns.

*LGeneral* contains little documentation, but the game is operated entirely by menus and point-and-click – so, even if you haven't played the original, it shouldn't take too long to work out what's going on. The game is played against an artificially intelligent CPU opponent only, but no less addictive for that. This is very polished clone of classic game, that sure to enthrall those with a passion for old-fashioned war games.

## PARTITION EDITOR

# Parted

■ VERSION 1.6.1 ■ WEB [www.gnu.org/software/parted](http://www.gnu.org/software/parted)

Not that long ago the open source world lacked any competent tools for managing PC partitions. Thankfully, GNU's *Parted* appeared to fill this niche, providing the ability to simply and reliably create, copy and resize FAT and *Ext2* partitions, and many distros now use *Parted* for partition management during installation.

The next generation of *Parted*, the version 1.6.x series, has recently become stable, and increases *Parted's* usefulness immeasurably. Portability has been improved, as it now supports Sun and MIPS partition tables – as well as DOS, BSD, and Mac tables. Support for *ReiserFS* has now been added –

thanks to Yury Umanets's *libreiserfs*, (<http://reiserfs.osdn.org.ua>). This library ships with new tools for managing Reiser partitions, and when *Parted* is built with *libreiserfs*, you can create, copy, and non-destructively resize Reiser filesystems.

The other new feature in the new generation of *Parted* is a rescue function. This is immensely useful if you lose your partition table and don't have a back-up. It works in the same manner as *GPart*, and, given a range of blocks to search, it can scan a drive block by block to detect any lost partitions there – and so rebuild your partition table. *GPart* has saved my life on a couple of occasions, so it's great to see this ability in *Parted*.

```

#kpartfs PART-TYPE FS-TYPE START END      make a partition with a filesystem
move MINOR START END      move partition MINOR
name MINOR NAME            name partition MINOR NAME
print [MINOR]             display the partition table, or a partition
quit                      exit program
rescue START END          rescue lost partitions between START and END
resize MINOR START END    resize filesystem on partition MINOR
rm MINOR                  delete partition MINOR
select DEVICE              choose the device to edit
set MINOR FLAG STATE      change a flag on partition MINOR
(parted)
(parted) print
Disk geometry for /dev/ide/host0/bus0/target0/lun0/disc: 0 008-19478 937 megabyte
es
Disk label type: msdos
Minor  Start      End          Type          Filesystem  Flags
1      0 031         956 995      primary       reiserfs
2      956 997      18386 894    extended
3      957 078       1280 168     logical       linux-swaps
4      1200 199       6918 618     logical       reiserfs     boot
5      6918 649       8824 768     logical       reiserfs
6      8824 799       16457 211    logical       reiserfs
7      16457 243     18087 526    logical       ext3
8      18386 895     19469 399    primary       ext3         boot
(parted) resize 3 18386 895 19000

```

It ain't pretty, but *Parted* – now with *ReiserFS* support – is a dependable tool for managing your disk partitions.

What's not so great about *Parted* is that the standard text-only frontend still has a hideous interface when used interactively. To confuse things further it has a separate partition naming scheme from the kernel's and *GRUB's*. You have to select a disk, and then address a partition by number. It

still has no progress reporting, either, so when performing long operations, you often wonder if it is ever going to finish. Still, it works, and it's certainly more pleasant than *FIPS*. I'm still waiting for somebody to produce a general-purpose, graphical interface to *Parted*.

## DATA MAPPING

# Generic Mapping Tools

■ **VERSION** 3.4.1 ■ **WEB** <http://gmt.soest.hawaii.edu>

One of the areas where the notion of free software clearly excels is in creating tools for science and academia. There is no point, from a scientist's point of view, in presenting data in a form only you can read. Open standards prevail, and evolve with the needs of the people who use them. Knowledge, and code is shared, and everyone can get one with doing what they need to without worrying about buying many different software applications. The *Generic Mapping Tools* suite clearly illustrates this approach.

The suite comprises of a range of tools designed primarily for the mapping of data – whatever form of data that might be.

The GMT tools have standardised on the netcdf file format for data storage, but include tools for converting other files to this format. Sorry, but we weren't able to include the NetCDF source on the discs this time, but you can obtain it and instructions on how to compile from the Unidata website [www.unidata.ucar.edu/packages/netcdf/INSTALL.html](http://www.unidata.ucar.edu/packages/netcdf/INSTALL.html)

In pretty much constant development, these tools provide an easy way of generating maps and histograms of all sorts of data. Most of the included applications generate postscript output, which is very portable and easy to manipulate.

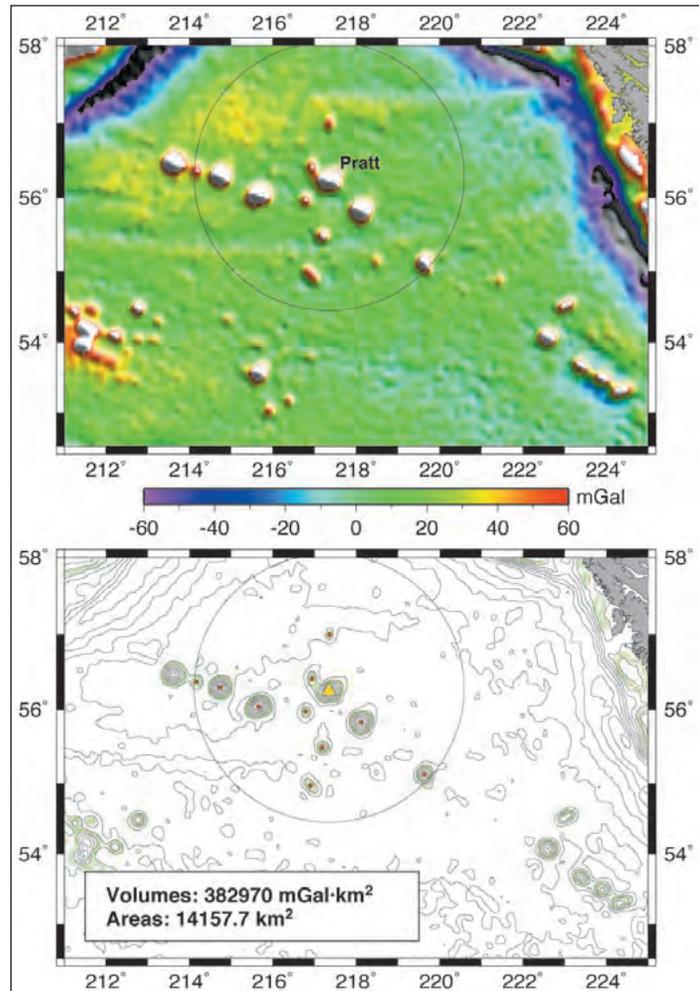
It certainly beats developing your own software to generate, for example,

a contour map from some digital elevation data. There are a total of 54 different commands included in the GMT collection, most of which are for manipulating gridded data points and creating a postscript output.

The command can be used on their own, but are more often used in sequence to generate detailed maps. For example, the **grdlandmask** command can be used to create a mask for data above or below the shoreline, from coastal data. A subsequent **grdcontour** can generate a contour map for the resultant grid file, so contour data is only displayed for the landmass. Various other tools are included to make the output more pleasant and useful – the **psbasemap** command for example, can draw a grid overlay with latitude and longitude marked on the axes, with optional tickmarks and other reference data.

As well as the tools themselves, the various GMT packages include example scripts to illustrate the type of data which can be produced – invaluable for understanding how to combine the output into a single file.

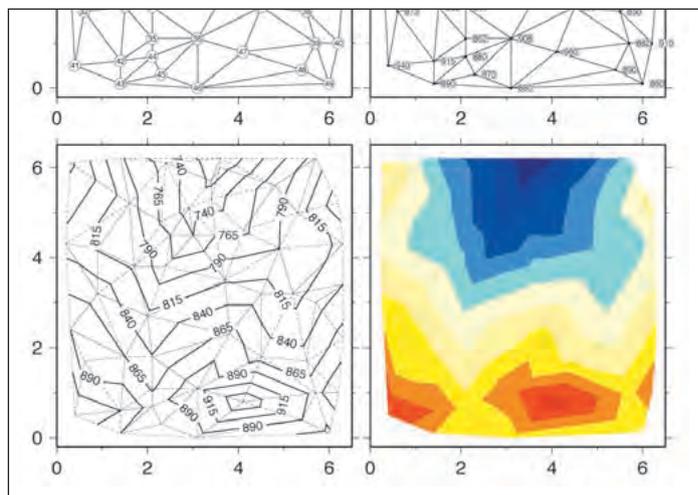
Man pages are provided, and there is some documentation, but it has to be said that these tools are very complicated, mainly due to the complexity of the tasks they can achieve. For example, there are various different mapping methods and projections – Mercator, Gnomonic, Orthographic, Cassin, etc. – and many



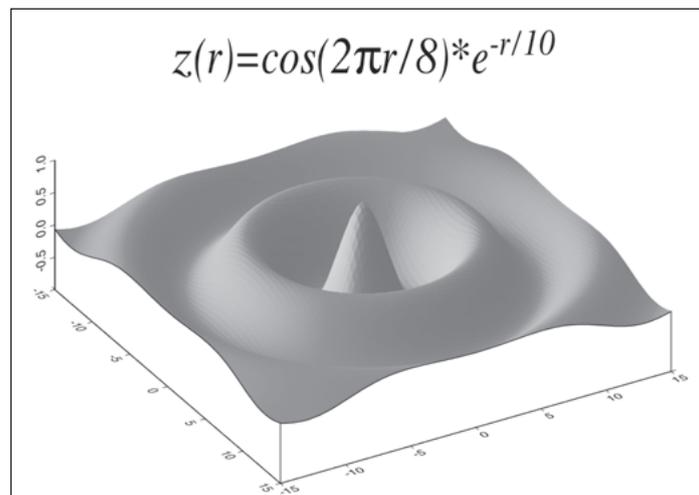
Various rendering methods, including relief maps, are supported

different ways the parameters can be specified. It will take some investment

in learning how to use these tools before you can really benefit. [LXF](#)



Combining several views in the same document is quite easy.

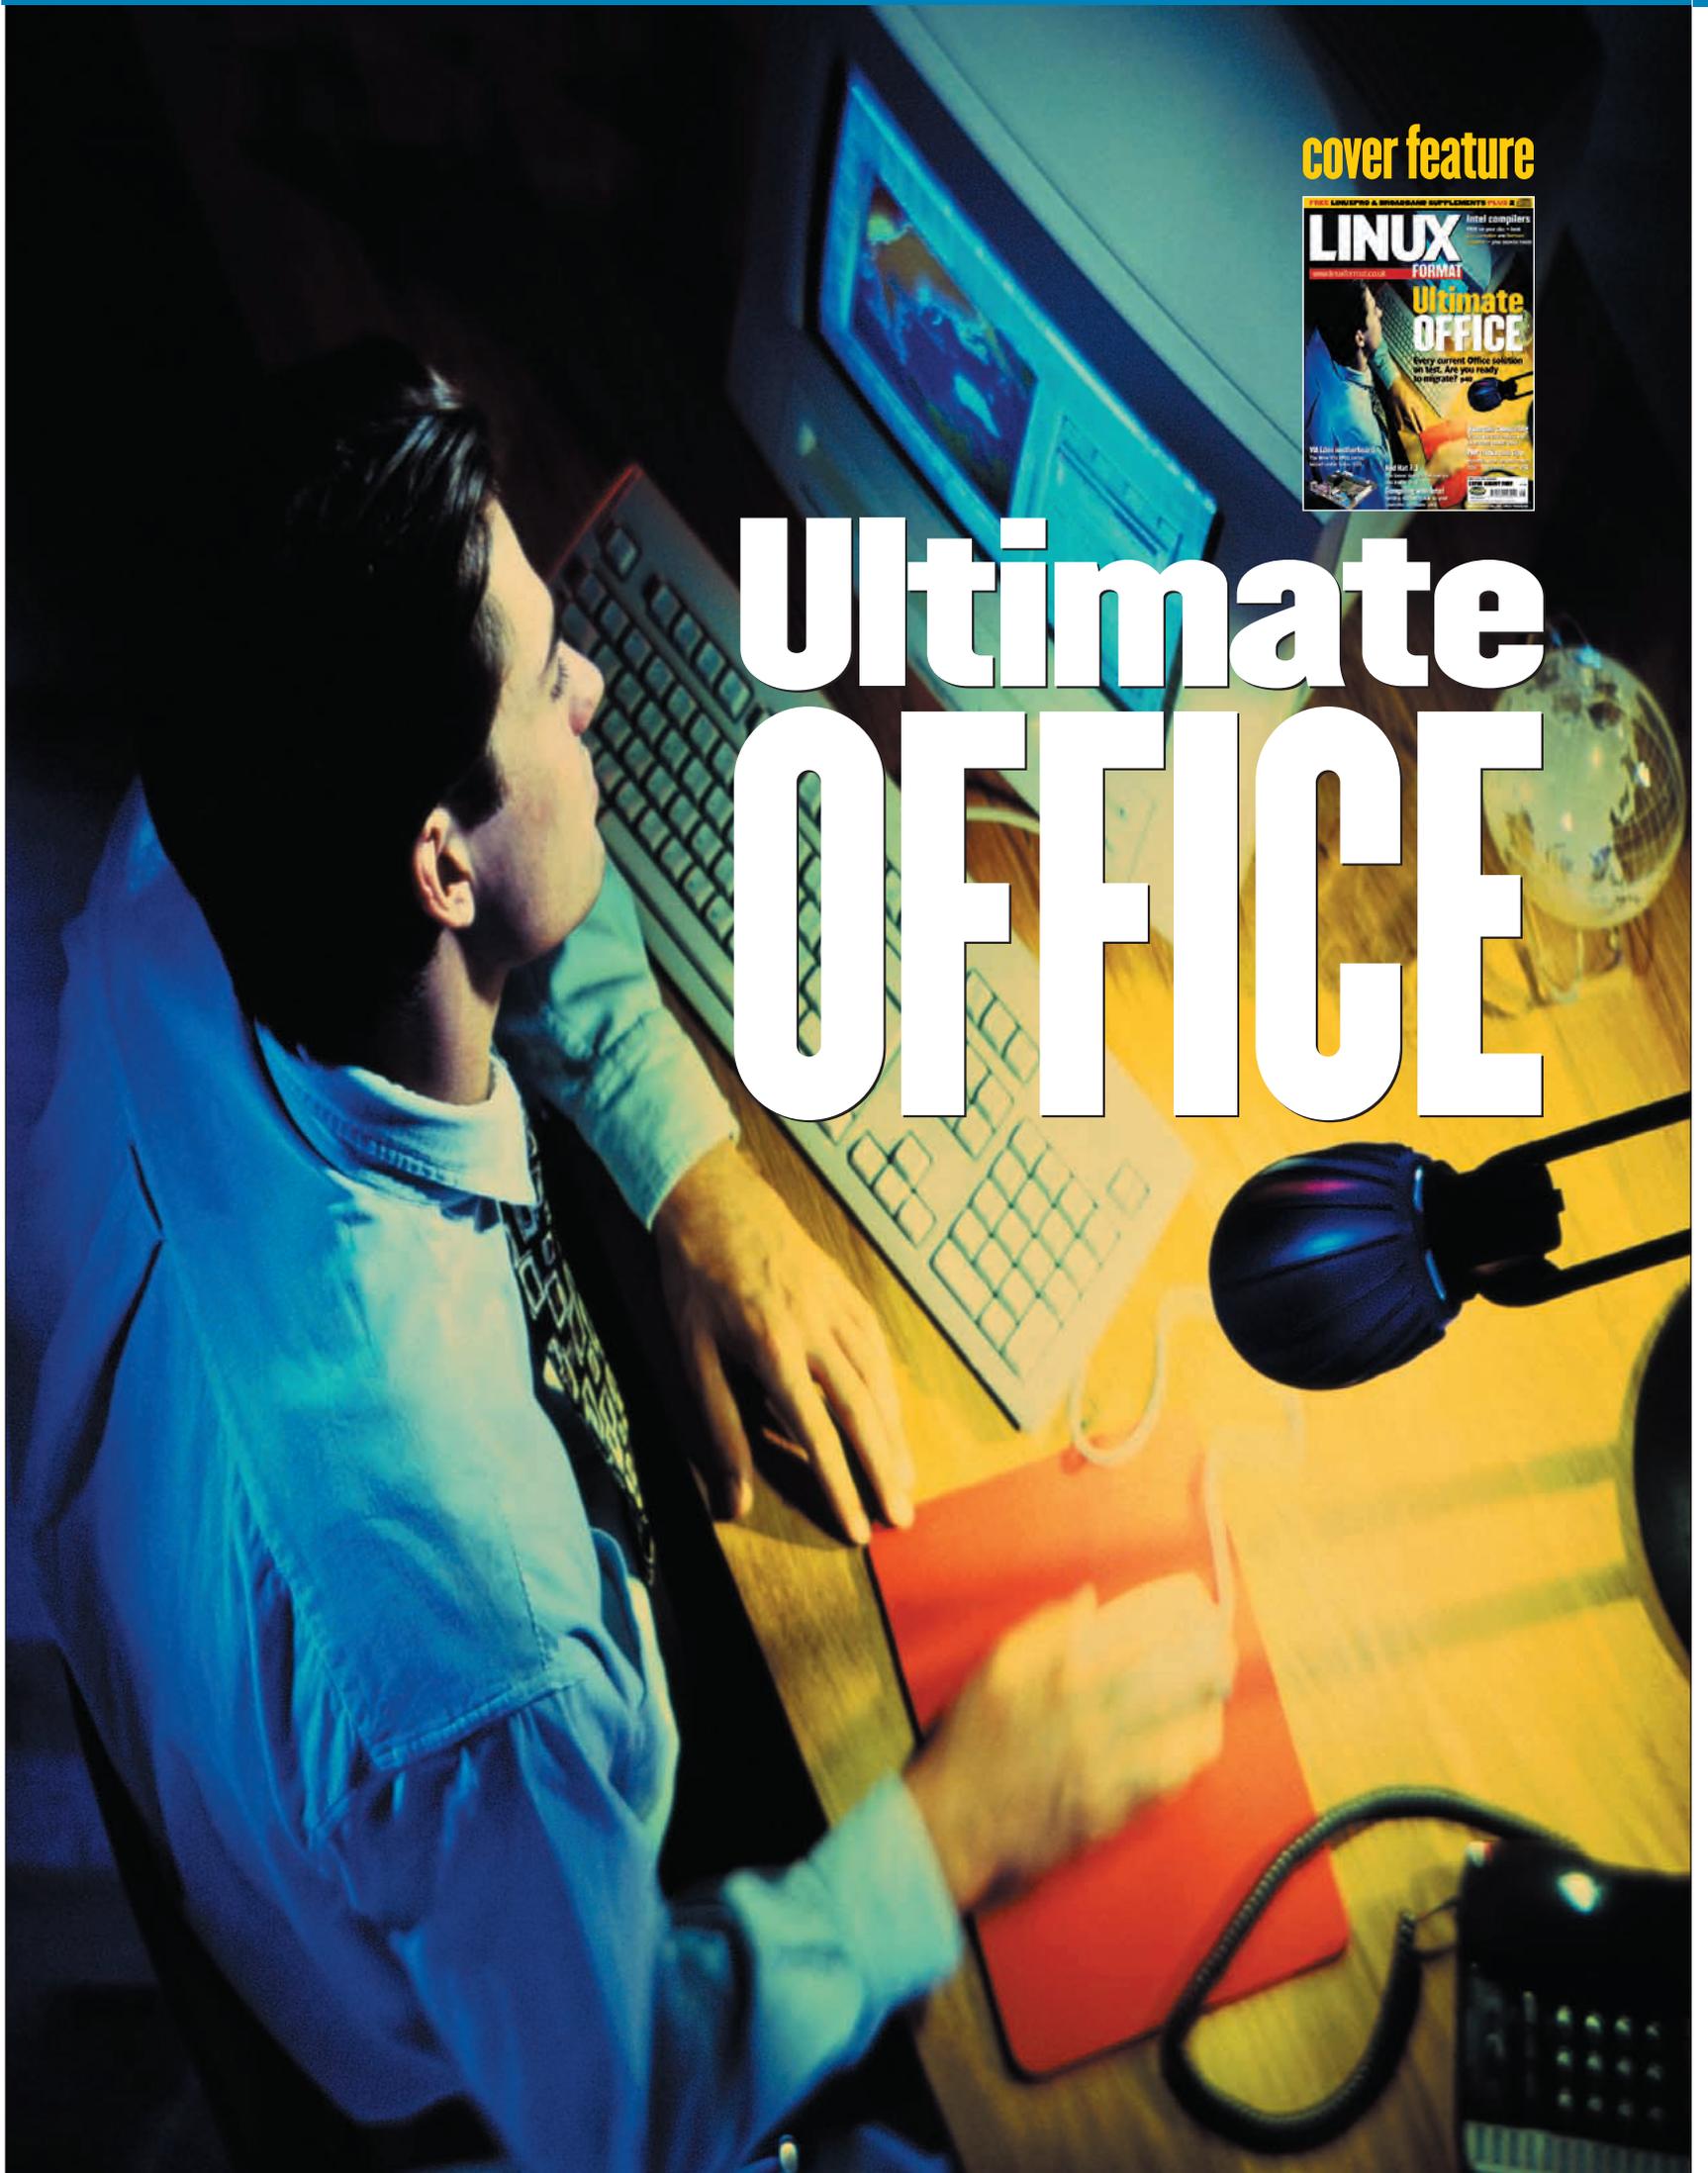


Your data might not be a real map - it's useful for other graphics too!

cover feature



# Ultimate OFFICE



## The biggest feature ever in *Linux Format*, Nick Veitch introduces the most comprehensive guide to Linux office software in the history of mankind.

Quite a few people seem to have a fixation with the “is Linux ready for the desktop” debate. I’m sure amongst our readership you would find strong proponents on either side of this argument. But what is certainly clear is that Linux is now well served in the area that more than 80% of us use desktop boxes for – office software. Whether it’s compiling sales reports, mailmerging form letters, sending faxes, writing docs, analysing data or 101 other common office tasks, it seems Linux is ready to meet this challenge.

But although the software exists, there is still the question of whether it actually works well enough to replace whatever systems you are using currently. Fortunately, these Linux apps are mostly inexpensive, so trying them out isn’t going to be costly, except in terms of time.

That’s where we come in. In the most comprehensive roundup of Linux Office software ever, we analyse each component of the six contenders in a variety of different categories. No less than four authors have sweated over image creation, presentation design,

complex spreadsheets and document creation to bring you what can only be described as the Ultimate Guide to Linux Office software.

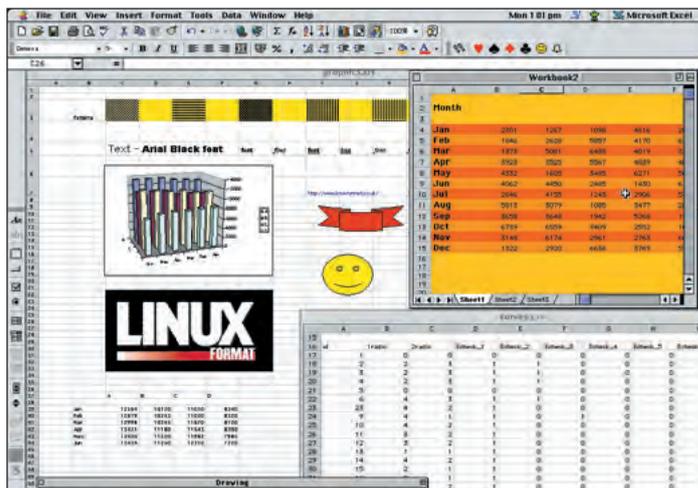
### Ease of Use

One of the primary concerns with this type of software is “will any monkey be able to use it”? An application that may work well for the Linux *cognoscenti*, may be inappropriate for deployment in a general office environment. Online help, clear interface design, easy to understand file requestors, menus and print options are the order of the day. You can be assured that whoever emerges the victor in this contest will have to pay close attention to ease of use.

### Compatibility

Of course, we have to live with the fact that perhaps other people might not share our choice of office software. Unlike some major software manufacturers, we will have to concede that other software exists, and some people may choose to use it. Compatibility in Office software is pretty important. If you are a small business running a single solution, you

**“Whilst integration is often a benefit, there’s a lot to be said for the Unix philosophy of picking individual tools”**



Will your favourite Linux app open those troublesome Windows files?

## Installation

Too easy to comment

Thankfully, this is one of those features we can do without worrying too much about people complaining that x, y, or z wouldn’t install. We haven’t commented on the installers in the main text, because they all work. The GNOME

Office and *KOffice* suites will probably be installed with your distro, and if not, there are certainly packages available for your distro. All the others come with trouble free installers, which should present no real difficulty.



may not be too worried about sharing data. If you have to deal with other companies, as a client or a customer, you will soon find that everyone expects you to be able to open the documents they send you in increasingly obscure formats (If we had a penny for every time a PR company sent us a .doc or .ppt file...).

Compatibility is almost certain to be an issue if you are considering migrating your existing systems to Linux, or from one Linux solution to another. The good news is, that amongst open source solutions, the move is towards freely exchangeable data in a standard format – XML. The huge advantage of XML as a format for storing styled documents or spreadsheet data is that the resulting files are easily readable by other systems (the key to file compatibility between, e.g., *KSpread* and *Gnumeric*).

Proprietary software may cling to deliberately obfuscated formats though, so we have tested compatibility thoroughly throughout all the components tested. For most, an ability to work well with Microsoft *Office* files is the prime requisite, so we have included as much detail as we could glean about the compatibility of Linux solutions in both reading and writing MS Office files.

### The contenders

These are the six latest, most comprehensive and popular office

suites of the moment. We have excluded some, like *SOT Office*, because essentially this is a rebranded version of *OpenOffice.org*.

As well as Open Office, we have included the readily available *KOffice* suite, and the loose confederation of GNOME Office tools. While integration is often a benefit (and we’ll see that most of the suites on test are good at this), there is a lot to be said for the Unix philosophy of picking individual tools to do specific jobs.

The GNOME Office tools are the HiFi separates of the office suite world, but the flexibility they offer could be very attractive. For example, you may chose to use *Hancorn Office* as a suite, but use *The GIMP* for image editing. This is ‘Free Software’ after all. Now, turn the page and let’s get started.

## Databases

Access all areas

Unlike other desktop operating systems, one common application Linux does lack is a standalone flat file database. The reason for this is quite simply, with the range of powerful database servers for Linux, most software integrates with the likes of *MySQL*, *PostgreSQL* and others for data storage. In this roundup, only *StarOffice 6.0* comes with it’s own database server (*adabas*), but will also integrate with other common database servers.

# Word processors

**Richard Drummond** puts six Linux word processors head to head, to find out which is the best one for writing your next novel, manual, fax or memo.

The most commonly-used component of the office suite will be without doubt the word processor. For many, it will be the strength of the word processor alone that will make or break the suite.

The features that you will look for in a word processor will vary greatly depending on your intended use. Someone who requires only to knock up the occasional letter will have very different requirements from somebody writing a PhD thesis. The packages on test here have something to offer to both. Of course, compatibility with MS products will be a deciding factor, especially in the work place, and all word processors we tested claim some ability to at least import *Word* docs.

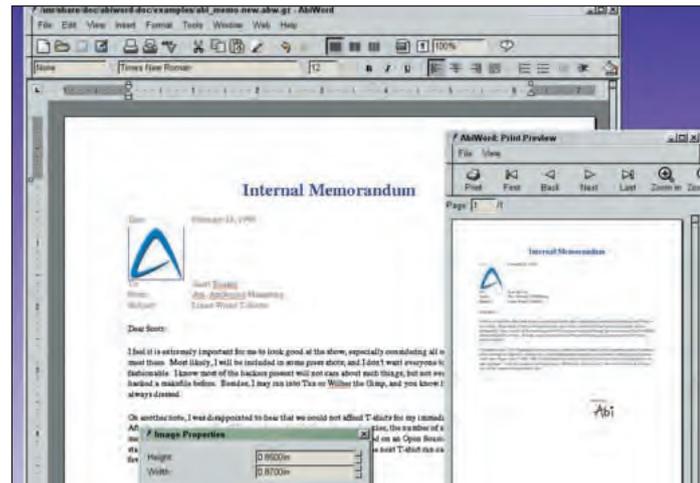
## ABIWORD

*AbiWord* is in many ways the odd man out of the word processors on test. Part of the loose confederation of GNOME Office packages, it really is a stand-alone product. Consequently, it's a lot smaller and lighter than the competition, with lower memory requirements and fast start up times.

*AbiWord's* simple and efficient user interface provides all the basic style and formatting controls you need. It lacks any kind of support for frames or tables, however, and so is less good at complex document layouts; also, it cannot embed spreadsheets, formulae or graphics, like the many of the other packages that are part of a suite can. It is able to insert images inline, and can handle any image format that the GNOME libraries can. Structured and vector formats are less well represented than pixmap formats, and so a lot of clipart will not be directly useable. A basic and useful set of document fields are provided – such as dates, word counts, page numbers – but no there's no mailmerge facility.

*AbiWord* offers a rich set of filters for importing and exporting documents to external file formats. Of course, MS *Word* import is supported, and this works surprisingly well. There are also filters for importing documents from *Applix (Anyware) Words*, *Word Perfect*, *Psion Word* and *RTF*.

Despite its limited functionality, one significant weapon in *AbiWord's*



***AbiWord* is a simple, light yet surprisingly functional word processor.**

arsenal is that it supports plug-ins. Default plug-ins include a thesaurus, plus plug-ins that let you query online dictionaries and encyclopedias or translate via *Babelfish*, and one that lets you edit inline images in the *Gimp*.

*AbiWord* is still a young project, so it's less advanced than most of the other packages on test here. It's small size and ease-of-use means that it's worth having around for simple jobs,

even if you have *OpenOffice.org* or *StarOffice* installed.

**AbiWord** **VERDICT**

A small and light word processor with a surprising amount of functionality and good import functions.

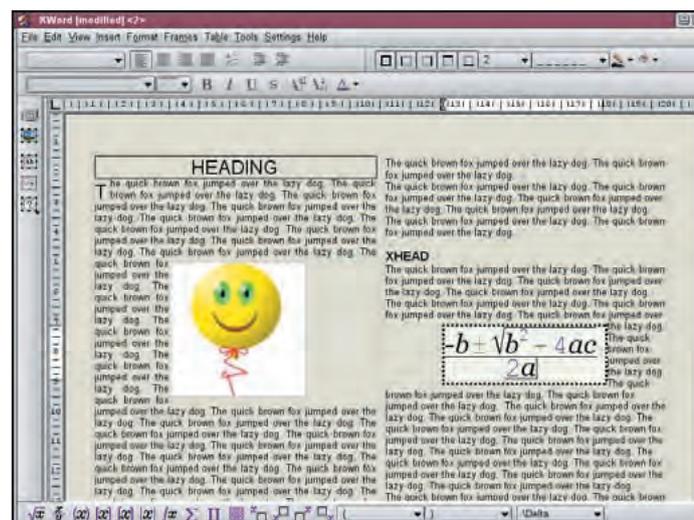
**LinuxFormat RATING**

██████████ 7/10

## KWORD

*KWord* shows a lot of promise, but is probably the weakest package on test. Its principal advantage is that it is based on KDE's component architecture and gives great integration between the packages that make up the *KOffice* suite. You can embed a *KSpread* spreadsheet in a *KWord* document (or vice versa) for example. You can also embed a *KWord* part in *Konqueror* to browse *KWord* docs on your desktop.

While it lacks a lot of the features of some of the more mature word processors, *KWord's* handling of frames give it the makings of an excellent document processor. There's still a lot of work to be done – you can't currently flow text around both sides of a frame – but otherwise it is good at quite complex layouts.



***KWord* is quite capable at lay-outs, but otherwise under-powered.**

*KWord* has a simple, uncluttered interface. One reason for this is that it is less functional than some of its

competitors, so there is less to burden the interface. However, KDE provides an excellent toolkit for building apps,

and this is made good use of in the *KOffice* suite. Selecting an embedded *KOffice* object on the page, integrates into the *KWord* window the interface of the application that handles that type of object.

Still, a lot about *KWord* is less exemplary. Stability is poor, and, what's worse, the autosave function doesn't seem to work. Its ability to import MS *Word* documents by far the weakest of those on test. However, *KOffice* is still a very young project, so its limited functionality is to be expected.

**KWord** **VERDICT**

Shows promise, but currently offers too little functionality for most needs.

**LinuxFormat RATING**

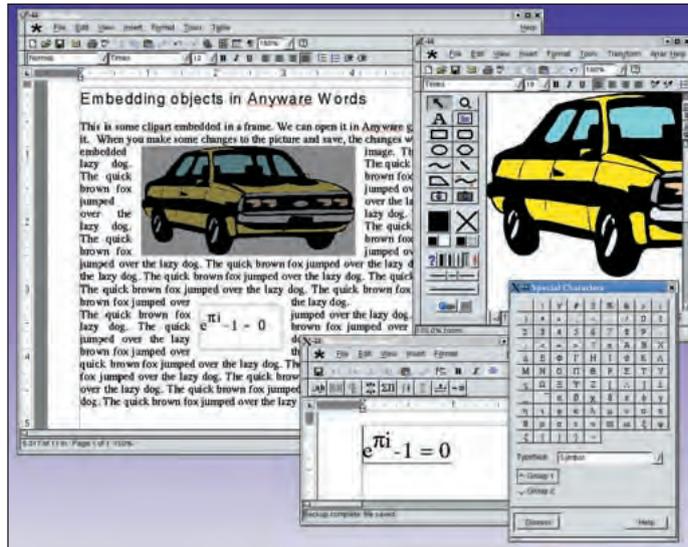
██████████ 5/10

## ANYWARE WORDS

*Anyware Office* (née *Applixware Office*) was the one package that I hadn't seen before doing this round-up, and it was the one that surprised me the most.

When you first load up *Anyware Words*, you are presented with a simple, clean *GTK+* interface – which belies the functionality hidden underneath. *Anyware Words'* interface layout is refreshingly different from the other packages on test, and it doesn't try to be yet another rip off of Microsoft *Word*. It may require a period of adjustment to get used to the different layout and terminology, but *Anyware Words* is no more difficult to use than the other packages on test. The package includes excellent online documentation which will help you find your way around.

The interface is a little cumbersome in some areas, however, particularly with embedding frames and objects. For example, unlike the other packages, *Anyware Words* doesn't make much use of context-sensitive menus. If you select an object, most operations on that object are available only via the main toolbar – which slows you down; and, if you



***Anyware Office* provides a simpler and cleaner way of working with embedded objects than *KWord* or *StarOffice* do.**

choose the 'Insert clipart' function to include a drawing from *Anyware Office's* bundled collection of images, then the image will be inserted inline by default – not in a frame. This is generally not what you want. You then have to create a frame, and cut-and-paste the image into it. Another complaint is that when you insert an image, its default size is set to be the size of the whole page – which again is generally not what you want.

Having said all this, I actually prefer the way that *Anyware Words* handles embedded objects compared to the other suites. With *KWord* or *OpenOffice.org/StarOffice*, when you insert an object, you can double-click that object to edit it. This will embed the GUI for editing that object right there on your word processor page. I find this approach rather cluttered and confusing. *Anyware Words*, on the other hand, pops up a new window for

you to edit that object. For example, if you insert a graphic, double-clicking it will bring up *Anyware Graphics* in a new window to edit that graphic. Do your modifications, click save, and the changes will appear in the image in *Anyware Words*. Simple.

*Anyware Office* includes a full complement of tools for most word-processing tasks: tables, lists, frames, formulae, a thesaurus, and cross-referencing tools for contents, indices and footnotes and endnotes. It has sophisticated typographical controls such as hyphenation and removal of widows and orphans. It also boasts a rudimentary yet functional mail-merge facility, which supports flat files only as a data source for merging. The icing on the cake is a powerful, built-in macro language and a rich set of filters for swapping documents with apps such as MS *Word*, MS *Write*, *Word Perfect* and *FrameMaker*.

### Anyware Words VERDICT

A great all rounder, *Anyware Words* combines ease-of-use with advanced features such as mail-merging and macros.

LinuxFormat RATING  

 9/10

## Fonts

### Your path to perfectly printed documents

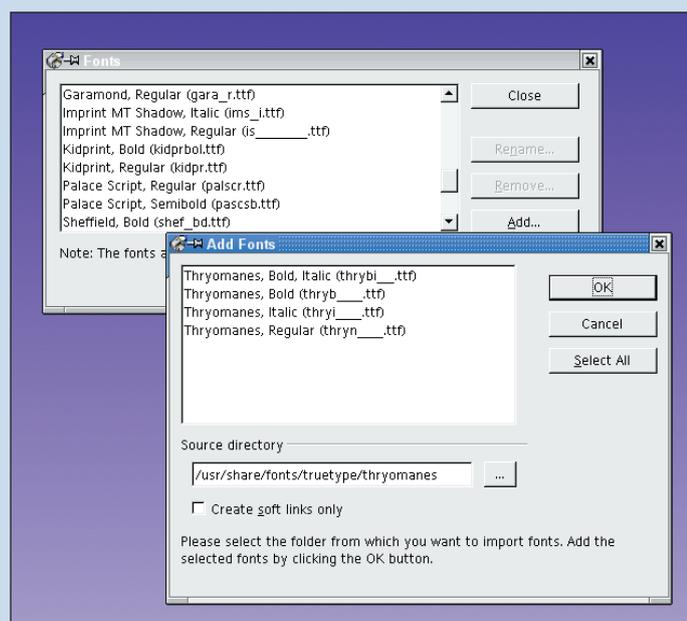
The handling of fonts causes special problems for office software under X, but it's perhaps more acute with word processors because the quality of hard copy is more important than with other office components. An application can quite happily use X's fonts for on-screen layout, but when it comes to printing, X doesn't provide accurate enough metric information to produce high quality output. Thus an application needs direct access to the font files themselves, and so needs to know where the fonts it uses are stored on disc. The various applications on test here solve this problem in one of two ways.

One solution is to by-pass X's font handling entirely. *Hancom Office*, for example, takes this route. It can only use fonts listed in its own personal font store: it cannot use X's fonts and (as far as we can tell – no documentation was available in the evaluation version) it only supports TrueType fonts.

The other solution is to use X fonts, but only those served from known file

paths, so the the fonts files are directly accessible. So, for example, *AbiWord* also has own store of fonts, but, in contrast, does use X's fonts handling and adds its own font path to your X server's font path when it starts up. *StarOffice* and *OpenOffice.org* behave similarly, but they also let you install fonts in a folder in your home directory and any fonts accessible to X. *Anyware Office* goes the whole hog and actually installs and uses its own custom font server.

Whichever solution an application chooses, it shouldn't really make much difference to the end user, except when it comes to making additional fonts available for use by the word processor. Both *Anyware* and *OpenOffice.org/StarOffice* come with tools for installing fonts, but with the others, you'll have to install any extra fonts manually. With *Hancom Office*, this is simply a matter of copying fonts into its font drawer, but with *AbiWord* and *KWord* you'll have to update the X fonts.dir and fonts.scale in drawers where you install fonts.



***OpenOffice.org* and *StarOffice* include a tool for installing new fonts in your local font directory.**

# OfficeWordProcessors

## HANCOM WORD

Despite ostensibly being part of the *Hancom Office* suite, *Hancom Word* feels less of an integrated component than the other word processors the constitute part of a suite. This is perhaps because *Hancom Word* pre-dates the other components in *Hancom Office*. Anyway, although all the applications in *Hancom Office* are based on Qt 3.0, *Hancom Word* has a different look and feel. You cannot embed components from the *Hancom Office* suite like you can with the other office packages on test.

On first loading up *Hancom Word*, my first reaction was of horror. Its interface looks vastly over-complicated, with four of five rows of floating toolbars, all densely packed with garish, little buttons. A little bit of

perseverance revealed that its interface isn't wildly different from the others – it's just that it's all displayed to you at once. They could do with making a better first impression. It's also worth noting, as well, that while *Hancom Office* use *libqt3* for its widget set, it uses its own copy of *Qt*, and so any *Qt* themes that you have installed (whether default *Qt* themes or those from KDE) won't be available to it – which is a pity.

*Hancom Word* hails from Korea, so, as you might expect, one area where it really scores is in support for internationalisation. Support for international fonts is good – in fact, font handling all round is impressive, with a wide range of effects and styles – such as outline, emboss, drop-caps, and over-typing – included. English, Chinese and Korean



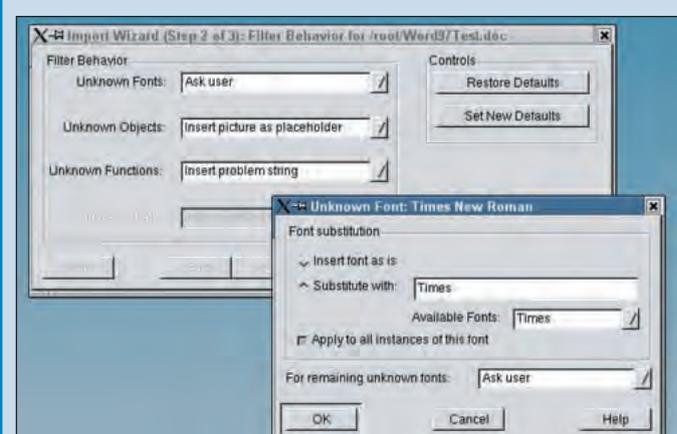
**Hancom Word has some great tools for producing font styles and effects.**

dictionaries were supplied with the version we tested, and the spell-checker is flexible, allowing you to apply all three dictionaries to a document simultaneously.

In other areas, *Hancom Word* is capable, but not exciting, and handling of tables, lists, pictures and frames all work as you might expect. The lack of support for inserting other objects

## Word compatibility

### Understandable attachments



#### Anywhere Words is good at importing MS Word documents.

*KWord's Word* import function is poor. It manages the basics and imports text, list and tables, but handling of page and paragraph formatting is woeful. It makes an attempt at frames, but can't cope with images or embedded objects. Fields are imported by value only.

*AbiWord* fares a lot better, especially considering it doesn't support tables or frames itself. It can, however, import textual content from both. Paragraph and page formatting are both good, and it supports some fields. Importing of pictures is poor, due to the lack of support for frames and WMF. Stability proved to be a problem while importing *Word* documents under testing.

*Hancom Word's MS Word* import function is adequate, but support for frames is poor, and, since fields are

unsupported, they are imported by value only. Stability is also a problem.

*Anywhere Words'* import function is excellent – almost on a par with *OpenOffice.org* and *StarOffice*. A novel feature is that it is an interactive process. *Anywhere* pops up a dialog asking you how to handle fonts not available to *Anywhere* and how to handle embedded objects and macros it doesn't understand.

*OpenOffice.org Writer* and *StarOffice Writer* are the most capable at importing *Word* documents, and showed no apparent problems with the features that we tested. Like *Anywhere Words*, they even attempt to convert and import embedded OLE objects such as spreadsheets and graphics. Export of *Word* documents is similarly complete.

## OPENOFFICE/STAR OFFICE WRITER

*OpenOffice.org Writer* and its commercial cousin *StarOffice Writer* are huge applications, but, as far as the *Writer* component is concerned, there's little difference between the two other than that *StarOffice* includes a thesaurus. *StarOffice* also has more documentation, fonts, clipart, sample documents – but, after all, that's what you're paying for. For now I'll refer to them both as just '*Writer*'.

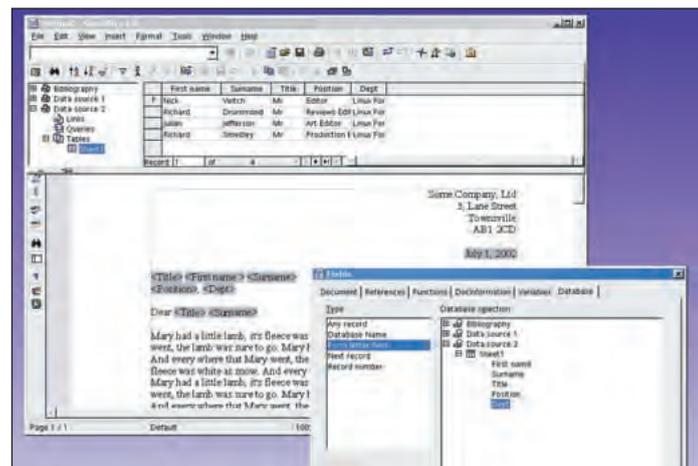
*Writer* is easily the most functional word processor on test here. Anything the others can do, it can do – if not better – than at least with more knobs on. It's an equal to Microsoft *Word* for most tasks. It even looks and feels a

lot like *Word*, so user's crossing over shouldn't face too much of a shock.

*OpenOffice.org* and *StarOffice* both suffer from slow start-up times and high memory requirements, but this improves, relatively, when you open more than one of suite's components. Like *KWord*, *Writer* is one integrated component of a complete office package, and, like *KWord*, you can embed components inside each other. Thus you may plug-in a spreadsheet object to include a spreadsheet in your *Writer* doc and edit it right there.

*Writer* is more than capable for handling large documents and complex layouts. Tools for creating contents pages, indices, footnotes, end notes and bibliographies are included.

One particularly powerful aspect



**Writer includes some powerful mail-merging tools.**

from the *Hancom* suite may prove restrictive for some, especially that there appears to be no simple way of importing data from spreadsheets. *Hancom Word* has limited cross-referencing tools. Footnotes and endnotes are supported, and apparently numbered captions for pictures – although this didn't appear to work during testing (the lack of documentation in the evaluation version caused many such difficulties). Additionally, fields, macros and mail-merging are all un-implemented.

### Hancom Word VERDICT

A competent word processor, but, apart from neat font effects, doesn't really stand out from the crowd.

LinuxFormat RATING  
 7/10

of *Writer* is its mail merge function. You can choose a variety of data sources for merging including ODBC or JDBC databases, an LDAP directory, a *Mozilla* or *Netscape* address book, a spreadsheet or even just a flat file. *Writer* will automatically determine an appropriate record structure for your data, and allow you to view it on screen. Inserting merge fields in your template document is then just a case of popping open a dialog and picking the list of fields that have been discovered in your data as you need them.

Many will complain that *OpenOffice.org* and *StarOffice* are just too large – and, of course, there are features here that most will never have a need for. But *Writer* is a serious word processor and is quite capable of handling the document creation needs of most businesses. Real faults with *Writer* are few, but I would like to see some filters for *WordPerfect* docs and a function to import a spreadsheet as a table rather than an object.

### StarOffice VERDICT

Easily the most powerful word processor on test and a capable substitute for Microsoft *Word* in the workplace.

LinuxFormat RATING  
 10/10

## Comparisons

	AbiWord	Applix	Hancom Word	KWord	Open Office	Star Office
<b>FONTS</b>						
Type1	yes	no(1)	no	yes	yes	yes
TrueType	yes	yes	yes	yes	yes	yes
Can use X fonts?	no	no	no	yes	yes	yes
Font replacement	no	yes	no	no	yes	yes
<b>OBJECTS</b>						
Fields/Variables	yes	yes	no	yes	yes	yes
Lists	yes	yes	yes	yes	yes	yes
Tables	no	yes	yes	yes	yes	yes
Frames	no	yes	yes	yes	yes	yes
Formulae	no	yes	no	yes	yes	yes
Built-in drawing tools	no	no	yes	no	yes	yes
<b>TOOLS</b>						
Autospell	yes	yes	yes	no	yes	yes
Hyphenation	no	yes	no	no	yes	yes
Autocorrection	no	no	yes	yes	yes	yes
User dictionaries	no	yes	yes	no	yes	yes
Thesaurus	yes	yes	no	no	no	yes
<b>CROSS-REFERENCING</b>						
Contents	no	yes	yes	no	yes	yes
Bibliography	no	no	no	no	yes	yes
Index	no	yes	yes	no	yes	yes
Footnote/Endnotes	no	yes	yes	yes	yes	yes
<b>INPUT FORMATS</b>						
MS Word 97/2000/XP	yes	yes	yes	yes	yes	yes
MS Word 95	yes	yes	no	no	yes	yes
Word Perfect 6/7/8/9	yes	yes	no	no	no	no
RTF	yes	yes	yes	yes	yes	yes
<b>OUTPUT FORMATS</b>						
Word 97/2000/XP	yes (2)	yes (2)	yes	no	yes	yes
Word 95	no	yes (2)	no (3)	no	yes	yes
Word Perfect 6/7/8/9	no	5/6 only	no	no	no	no
HTML	yes	yes	yes	yes	yes	yes
RTF	yes	yes	yes	yes	yes	yes
<b>MISC</b>						
Forms	no	yes	no	no	yes	yes
Macros/scripting	no	yes	no	no	yes	yes
Mail merge	no	yes	no	no	yes	yes
File format	XML	Proprietary	Proprietary	XML	XML	XML
<b>WORD COMPATABILITY</b>						
Fonts, styles, effects	Good	Fair	Good	Fair	Good	Good
Paragraph formatting	Good	Good	Good	Poor	Good	Good
Lists	Fair	Good	Poor	Fair	Good	Good
Tables	Poor (4)	Good	Fair	Poor	Good	Good
Pictures	Poor	Good	Fair	Poor	Good	Good
Frames	Poor	Fair	Poor	Poor	Good	Good
Fields	Fair	Fair	Poor (5)	Poor (5)	Good	Good

#### Footnotes

1 Anyware installs its own font server and has a tool for importing fonts which should support Type1 fonts, but this did not work under testing.

2 Actually exports Word documents as RTF.

3 Word95 documents displayed as a blank page under testing.

4 Tables are unsupported, but table content is imported.

5 Field values only are imported.

## CONCLUSION

Which is the best word processor for you? Well, it depends on your needs. Examine the table of features above, find out which one does what you require. In the case of *KWord*, *OpenOffice.org* and *AbiWord* – which are free – download them, and take

them out for a spin to see if they suit.

All of the packages on test here have merit in their own ways. Even *KWord*, which is the least capable of the lot, is useful for those with light needs. I wouldn't want to start writing a novel on either *KWord* or *AbiWord*, but both packages are good enough for editing letters, articles, and so on.

*Hancom Word* and, particularly, *Anyware Words* are both good, but if you have serious needs of your word processor, it would be hard not to recommend *StarOffice Writer* (or, for those on a budget, *OpenOffice*). It's easy to use and boasts the most complete functionality – in fact, it is more than equal to Microsoft *Word*.

# Spreadsheets

Totting up the totals and filling up the cells, **Nick Veitch** analyses how the contenders fare with the second most popular office application.

Rather like word processors, spreadsheets can and are used for a huge range of tasks, from managing stock portfolios, tracking sales and stock levels to analysing experimental data – in other words for pretty much everything that revolves around numbers.

Given the range of tasks that they are used for, combining all the user desired features into one piece of software is therefore, a bit of a challenge.

The other, separate challenge is that Microsoft's *Excel* has become the *lingua franca* of spreadsheets, so for a lot of people, a strong degree of

compatibility is desirable (as indeed was the case when *Excel* had to perform similar feats to surpass the market leader *Lotus 1-2-3* many powercycles ago).

However, the Linux platform is blessed with the commercial offerings and a few open source contenders in this area. We are also including the

*Gnumeric* spreadsheet here, part of the loose confederation of GNOME Office applications.

We tested all of the spreadsheets on the range of functions, speed, reliability and ability to deal with alien data formats, with a special emphasis on importing and exporting *Excel* data.

## GNUMERIC

*Gnumeric* is part of the loose confederation of GNOME Office tools, and ably fills the role of spreadsheet. The goal of the developers here is nothing short of creating the best all-round spreadsheet possible, with a huge amount of development work.

Using the *GTK* interface, *Gnumeric* is quick to load and pleasant to look at. The colourful toolbars may be a bit over the top for restrained business readers, but nevertheless are well organised – the only criticism is that the two main toolbars and the menu bar take up a fair amount of space.

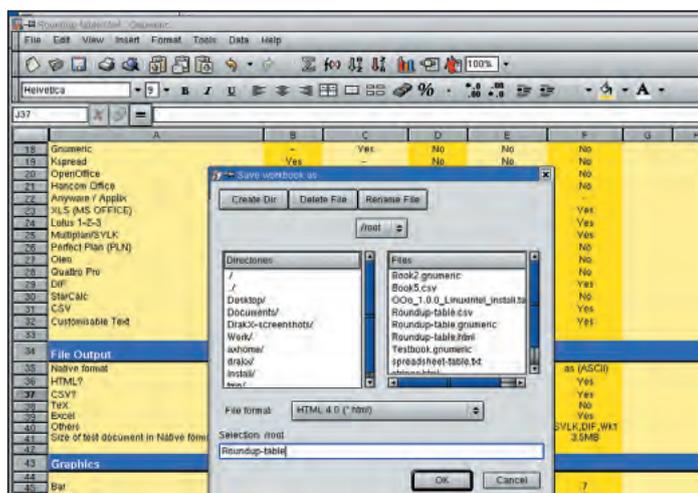
When it comes to functions *Gnumeric* has the edge over many of the offerings here. We could find less than 10 *Excel* functions that failed to work as expected. A great amount of effort has gone in to matching the *Excel* function set, and making it work

predictably (*i.e.* the same results, and the same aberrant behaviour when faulty data is included).

We tested the stable versions of all the software present, on the grounds that you wouldn't deploy unstable software for a mission critical task (we hope not anyway), but the current development version has gone a long way to fine tuning this compatibility.

As yet there is no chance of importing XL spreadsheets with Visual Basic macros. This too may become possible in the future though, as *GNOME Basic* is an ongoing project.

Graphics support is pretty poor in *Gnumeric*. It can include *Dia* images, but there is no bitmap support and no support for other graphics format. Even though it has a number of charting options, the range isn't great and it's hard to generate your charts externally. Full marks though for cunning – *Gnumeric* will convert *Excel*



**Gnumeric sports the greatest variety of output formats, including various shades of HTML document.**

Charts to one of its own formats if it can (*e.g.* a non-3D version).

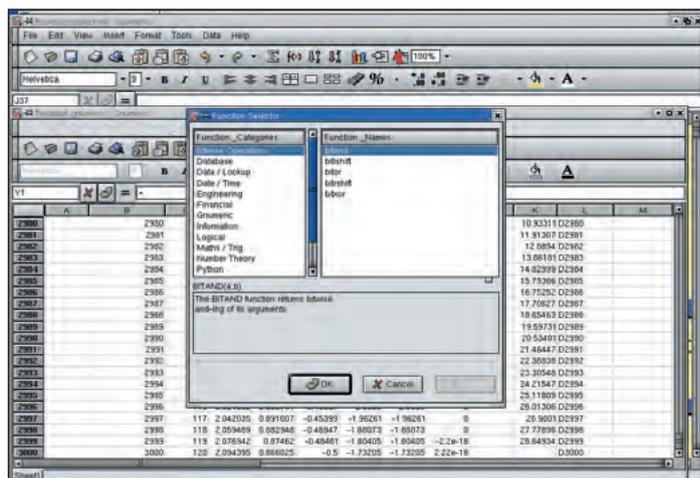
While relatively quick to load, it slows down terribly when dealing with very large files. Perhaps because of internal data organisation, it is very slow when performing operations like deleting rows of data from a large spreadsheet. Saving can also be slow.

A certain amount of extensibility is provided by the plug-in interface. This is responsible for providing support for most of the file format support, and also adds a python functions interface.

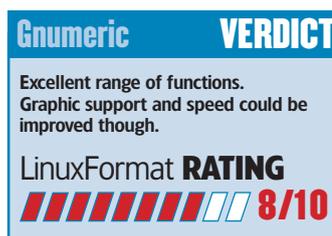
## KSPREAD (KOFFICE)

*KSpread* is the spreadsheet component of the *KOffice* suite, the office solution integrated with the KDE desktop. On the face of it, it should have a lot going for it. KDE is currently the most popular desktop environment, and the *KOffice* suite can take full advantage of the things that have made KDE popular – a good user interface and tight integration with other KDE apps.

Certainly it is easy to use, and easy to find your way around. The printing facilities are good and integration with other software, like spellcheckers, and support for embedded objects is good. It is even possible to create and edit embedded objects, such as a



**Gnumeric has the largest selection of built-in/plug-in functions.**



## STARCALC(STAROFFICE/OPEN OFFICE)

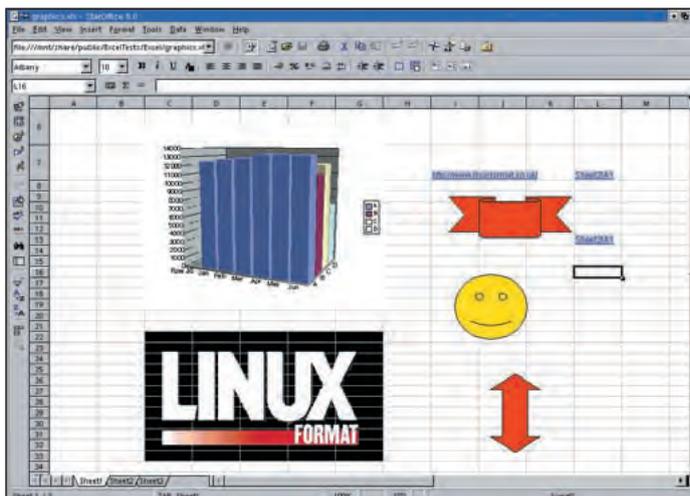
In terms of the spreadsheet function *StarOffice* and *OpenOffice.org* are identical, so are considered together.

Using a suspiciously *Motif*-style gadget set, the user interface is functional rather than attractive. Menus are a little clunky, but fully configurable. The toolbars don't take up too much screen area, but on the other hand, some of the icons are hard to understand (just as well the tooltips work okay).

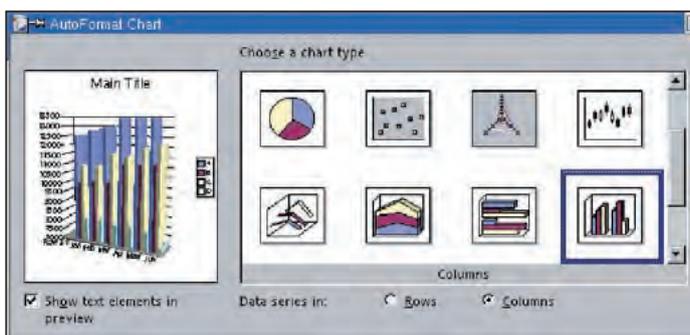
When it comes to features, there seems to be pretty much everything you could want. Conditional formatting, consolidation, goal seeking, various data tools and a fairly comprehensive Macro language. A large range of graphics formats are supported, including a good selection of vector/structured formats.

One of the major features of *StarCalc* is its ability to work with *Excel* documents. With almost 400 native functions, it's no surprise that all of the *Excel* basic functions are matched (although some are not quite so tolerant of foolish parameters, which may be worth watching).

This was the only spreadsheet to correctly represent all the graphics on our test *Excel* document, including proper transparency on the bitmap, an exact replica of the chart and correct



StarCalc's ability to read *Excel* spreadsheets is unsurpassed



Charting and graphics features are also very strong.

rendering of the *Excel* 'Autoshapes'. Even the links to other sheets and websites functioned properly.

There is also very limited support for XL macros. For example, elements of *Visual Basic* are syntactically the

same as *OpenOffice.org*'s (like the control of requestor windows for example) so simple macros may still work. Almost everything requiring calculations or manipulation of ranges of cell data will have to be rewritten

though. One bonus of *SO*'s approach to the macros is that the original *Excel* macros are saved – if you save the document out again in *Excel* format, the original macros should still work. That means you can make changes to *Excel* docs without removing any functionality.

When it comes to speed, both *SO* and *OpenOffice.org* (*OOo*) still suffer from slow startup times, although this has been improved. The times taken in the speed tests show improvements in several areas from previous versions, and while its native format might not be as quick as some, file sizes are very small (a 1.5MB *Excel* document was saved to a 200k file).

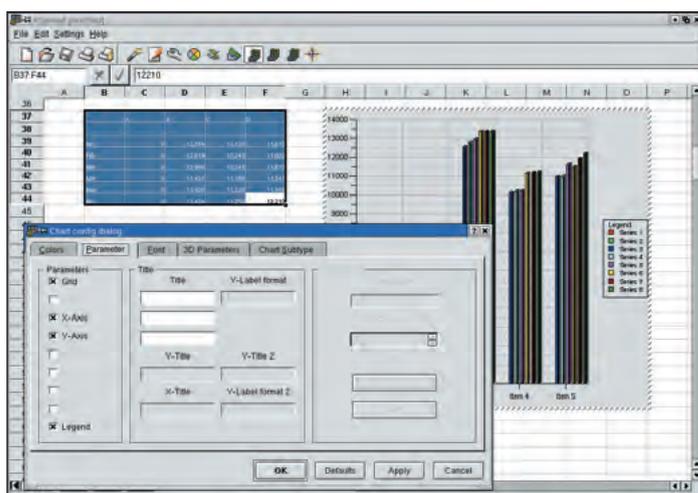
While functionally *OOo* and *SO* are very similar, there is some difference in the supplied material. *SO* includes better online help, and a large selection of interesting templates, including a useful Stock Portfolio manager which can update itself over the Internet! This, and extra help, can make getting to know the features of the spreadsheet that much easier.

**StarCalc** **VERDICT**

Hard to beat if you are after MS Office compatibility, and all-round solid performer.

**LinuxFormat RATING**

9/10



The integration with *KChart* is the high point of *KSpread*

drawing or flowchart, without leaving *KSpread*.

Where it tends to lose its way is simply in terms of being a fully

capable spreadsheet. There are a smattering of functions in the common groups, but nowhere near the comprehensive level offered by

the other contenders here (*Gnumeric* for example has at least three times as many functions).

This limitation means its effectiveness at importing *Excel* spreadsheets is fairly limited. In fact, the inclusion of functions which *KSpread* doesn't understand seems to result in the severe mangling of data. If you only want to import simple columns of figures, with simple arithmetic functions, *KSpread* will do a good job, but otherwise you'll be better off looking elsewhere.

*KSpread* also won't even try to convert *Excel* charts, which is a shame because *KChart* (which is a separate app used for charts and graphs) is very capable and easily able to create 3D charts of similar type.

Another area of concern is stability. On several occasions during testing,

*KSpread* crashed while trying to import documents in formats it claims to support.

The documentation with *KSpread* is also suitably nonexistent (the chapter headings are there, but no content) which makes investigating some of the features difficult. Overall it is fairly quick, lightweight and easy to use, but doesn't have the breadth and depth of other solutions. It might be fine for plotting a few sales figures, but for serious spreadsheet work, try one of the others.

**KSpread** **VERDICT**

User interface is good, good integration, but compatibility and function scope is poor.

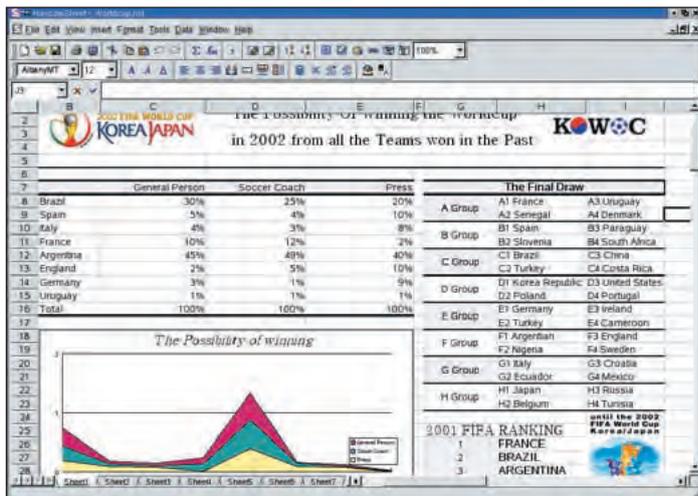
**LinuxFormat RATING**

5/10

# OfficeSpreadSheets

## HANCOMSHEET (HANCOM OFFICE)

The *HanCom* spreadsheet application looks pretty good. It has an uncomplicated interface, simple menu structure, a nice selection of fonts and some pretty neat features, like



**HancomSheet is colourful and easy to use.**

conditional cell formatting, a great data table wizard and a good range of functions.

Another nice feature is the ability to define custom lists for frequently used entries. There are pre-defined lists for days of the week, months of the year and so on, which can be

used to quickly create common labels (just enter the first item, choose a range and use Autofill).

Although falling short of the total number of functions offered by others like gnumeric, *HanComSheet* does manage to a lot of the common Excel functions, so from that point of view, you should be able to load most spreadsheets without too much of a problem.

For functions it can't understand it will insert the last known value as a constant, so most calculations should continue to work (it would be nice to be told which cells were affected though!)

An unfortunate problem with *Excel* support is that it doesn't seem to translate dates properly which could be rather a problem. *Excel* uses Microsoft's own datestamp format which is totally different from everyone else's. Dates imported into *HanComSheet* become, just numbers. If you style them up as dates, they will become dates, but completely wrong

ones, as *Hancom* seems to use a Unix datestamp for the native time format.

In terms of its graphics abilities, *Hancom Office* fares quite well. It copes well with *Excel* charts, though it makes a mess of autoshapes. Bitmap graphic support is also good, though there is no support for vector graphics other than its own charts.

Almost uniquely, *Hancom Office* won't load any filetypes other than its own, *Excel* files and CSV, or save them either.

The documentation is being almost constantly updated, and new PDFs can be downloaded from the Hancom website (you'll need to register first though).

**HancomSheet VERDICT**

Nice interface, some good features, fair *Excel* support. Could use better graphics support.

LinuxFormat **RATING**

/////// 7/10

## ANYWARE SPREADSHEET (ANYWARE OFFICE)

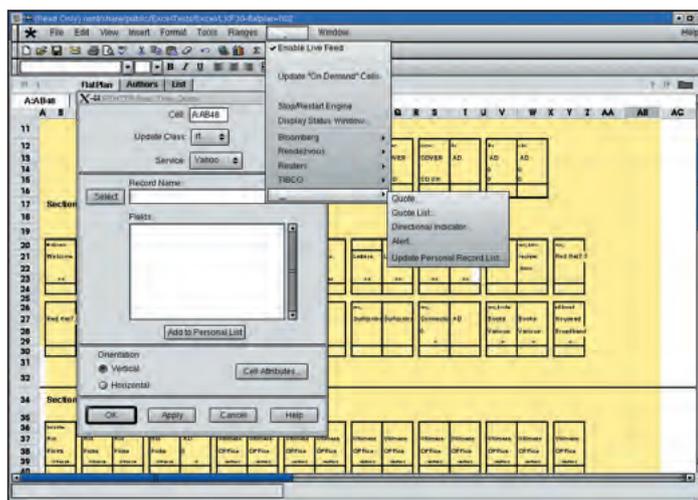
The java based *Anyware* spreadsheet is friendly but not overly cute, and as fast as you could ask for. The menus and toolbars are customisable to an extent, and unlike some of the other software considered here, manage to be clear and recognisable without taking up a huge amount of space on the screen.

When it comes to importing *Excel* documents, the *Anyware* spreadsheet goes about things in a sensible way. For a start, it lets you choose what to do if it encounters different types of errors (such as a formula not being recognised, or a font not being available. Secondly, it actually displays errors as the sheet is being loaded and will optionally create a log file. This almost makes up for the some of the deficiencies on the function front. A fair number of the *Excel* functions commonly used are supported, but there are some notable exceptions for more exotic functions, and some quite ordinary ones. The invaluable **COUNTIF** is missing as are

**BESSEL** functions, imaginary number support, format conversions (**Oct2Dec**, etc), and a few financial functions. In some cases similar functions exist, but use different parameters and/or names. This is a surprise as overall the *Excel* import function works well.

Although the software has support for hyperlinks (the interface for creating them looks very familiar to *Excel* users), it won't translate links from *Excel* sheets. You can create your own links to web pages and other files, but we couldn't find any way of altering the default behaviours (e.g. clicking on a web page opens that element in the *Anyware* HTML editor, rather than a browser). The links are also created without a label, so you will have to create one separately.

*Anyware* also make a fair stab at importing *Excel* graphics, including the Autoshapes – though the colours get a bit messed up. Charts are converted into *Anyware* style charts, which leave a little to be desired. A wide range of chart types are available, but the 3D options are less than inspiring – you get a fixed perspective and few customisation options. Although there



**The Anyware spreadsheet includes the capability to connect to their own RealTime data sources, for automatically updating documents.**

are a wide range of chart types, including the elusive Japanese candlestick charts (vital for tracking share prices), charting is definitely the weakest part of this spreadsheet.

Graphics support overall is very good though. Graphics are loaded though the *Anyware Graphics* package which supports a wide range of bitmap and vector file formats. Creating new graphics in a selected area opens up the *Anyware Graphics* component automatically.

One of the more unique features of *Anyware* is support for real-time data services (which requires further purchases). As a final note, it also supports *CUPS* printers.

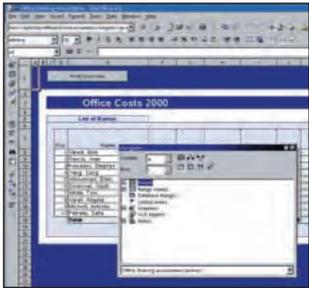
**Anywhere VERDICT**

Competent and capable, but not as *Excel* compatible as it may think

LinuxFormat **RATING**

/////// 8/10

## Conclusion



**Overall, the winner in this category has to be the StarOffice/OpenOffice.org duo**

For all round features, ease of use and compatibility the clear winner in this section is the *StarOffice/OpenOffice.org* software. If complete *Excel* compatibility is important to you, this is the best choice, although the *Anyware Office* does a good job in this regard. All of the spreadsheets (with the possible exception of *KSpread*) do a reasonable job of importing simple *Excel* documents though, and you could always save them out in some more universal format like CSV. The key here is how complicated the documents are in the first place – simple columns of numbers aren't going to present a problem for any of these packages.

As we said at the beginning, there are many uses for spreadsheets. If your use doesn't go much beyond tables of figures and the occasional date, any of the software discussed here would be fine. For numbers of functions, *Gnumeric* and *StarOffice* are good, and both provide the facility to custom define your own functions, although this will require some programming skill.

*Anyware Office* is also a good all-round performer, and while it isn't 100% *Excel* compatible, at least it tells you what incompatibilities it has found when loading documents. It also has some unique features which may be more suitable for groupworking environments.

*KSpread* doesn't handle large documents well (it failed to load the test large document from CSV, and crashed when we tried to generate a similarly sized document from within the software). It doesn't have a great scope of functions, or support for many file formats. On the plus side though, the charts are a lot more attractive than those generated by *Gnumeric* or *Anyware*, and it integrates well into the KDE environment. If you wanted to create graphs for presentations, etc, this is a good choice.

## Comparisons

Software	Gnumeric	Koffice (Kspread)	OpenOffice.org/StarOffice	Hancom Office	Anywhere Office
Version	1.0.4	1.1.1	1.0/6.0	2.0.1	
Interface	GTK	KDE 2	custom	custom	custom
<b>FORMULAE</b>					
Database	Yes	minimal	Yes	Yes	Yes
Engineering	Yes	minimal	Yes	No	minimal
Logical	Yes	Yes	Yes	Yes	Yes
Statistics	Yes	some	Yes	Yes	Yes
String	Yes	some	Yes	Yes	Yes
Complex numbers	Yes	Yes	Yes	No	No
<b>FILE INPUT</b>					
Gnumeric	-	Yes	No	No	No
Kspread	Yes	-	No	No	No
OpenOffice.org	No	No	-	No	No
Hancom Office	No	No	No	-	No
Anyware / Applix	Yes	Yes	No	No	-
XLS (MS OFFICE)	Yes	Yes	Yes	Yes	Yes
Lotus 1-2-3	Yes	No	Yes	No	Yes
Multiplan/SYLK	Yes	No	Yes	No	Yes
Perfect Plan (PLN)	Yes	No	No	No	No
Oleo	Yes	No	No	No	No
Quattro Pro	No *1	yes *3	No	No	No
DIF	Yes	No	Yes	No	Yes
StarCalc	No	No	Yes	No	No
CSV	Yes	Yes	Yes	Yes	Yes
Customisable Text	Yes	No	Yes	Yes	Yes
<b>FILE OUTPUT</b>					
Native format	Gnumeric (XML)	KDE (XML)	sxc (XML)	.Hst	.as (ASCII)
HTML?	Yes*2	Yes	Yes	Yes	Yes
CSV?	Yes	Yes	Yes	No	Yes
TeX	Yes	No	No	No	No
Excel	Yes	No	Yes	Yes	Yes
Others	DIF	Gnumeric	SYLK,DIF,dBASE	-	SYLK,DIF,Wk1
Size of test document in Native format	462k	-	193k	1.5MB	3.5MB
<b>GRAPHICS</b>					
Bar	3	3	5	6	7
Line	2	3	10	4	9
Pie	2	1	4	3	6
Scatter	1	0	4	2	10
Bubble	1	0	0	0	1
Area	0	3	3	5	6
3D Charts?	No	Yes	Yes	Yes	Yes
Legends/Labels?	Yes	Yes	Yes	Yes	Yes
Embedded objects	Yes	Yes	Yes	Yes	Yes
Bitmap support	No	No	Yes	Yes	Yes
Structured graphics	Dia	Kontour	StarOffice Draw	No	Various
<b>SPEED</b>					
Speed test 1 (open large doc.)	120s	Failed	23s	6s	7s
Speed test2 (delete row in large doc.)	13s	Failed	1s	8s	1s
test3 (save large doc.)	46s	-	27s	9s	6s
<b>MS EXCEL COMPATIBILITY</b>					
Dates	Yes	Poor	Yes	No	Yes
Links	No	No	Excellent	some	No
Graphics	Poor	Poor	Excellent	Good *4	Good*4
Styles	Fair	Fair	V.Good	Fair	Good
Charts	Fair	No	Excellent	V.Good	Fair
Macros	No	No	Some	No	No
<b>MISCELLANEOUS</b>					
User defined functions	via python API	No	Via Addin API	No	No
Macros	No	No	Yes	Yes	Yes

\*1 Development version includes this.

\*2 including HTML4.0 and HTML fragments.

\*3 Claimed support but test files didn't work. \*4 Missed transparency settings

# Graphics packages

Nobody buys an office suite solely for its graphics package, says **Andy Channelle**, but some of the design apps with the Linux office suites are strong selling points in their own right.

Once the word processor, spreadsheet and presentations are nailed, a developer's thoughts turn to the 'second division' apps; it may sound harsh, but no one uses an office suite for its design tools.

Adding a little visual flair to a headline or website menu, a fancy border to a text frame or editing the occasional piece of clip art are the extent to which these applications will be tested, so the key is not to do everything, but to do the basics well. And as suites tend to be operated by casual users, making the environment familiar is also desirable – of course KDE and GNOME offer 'standard' ways of doing things but integration means more than just having a unified key binding for Save As.

## STAROFFICE 6

*StarOffice*, and its open source counterpart *OpenOffice.org*, excel in the area of integration; double click on a text frame or recently drawn box, and the toolbar at the top of the window switches to include all the standard word processing tools. Similarly, if you're writing your thesis and discover you need to add a diagram, clicking the 'show draw functions' tool on the left of the screen brings up a small subset of the drawing tools from the illustration element. The other advantage of this is that you have access to all the usual proofing tools in the illustration module, so there's no excuse for poor spelling in any *StarOffice* project.

Many of the beta reviews saw the inclusion of a drawing package as an anachronism, citing the example of market leader Microsoft which doesn't include one in its flagship suite. I would disagree, especially when it comes to running the software on Linux which is not overly blessed with high quality drawing software.

*StarDraw* may not be able to compete feature for feature with products from Corel, Adobe or Macromedia, but for

everyday business graphics – editing clipart for a presentation, creating basic buttons for a website – it has everything you need.

A vertical bar on the left of the main window provides access to most of the basic drawing tools. Clicking and holding almost any of the buttons will launch a flyout with a subset of options; for instance, the rectangle box has options for square, rectangle, rounded corner square and rectangle, plus filled versions of each, while the curve button has options for bezier curves, polygons and freehand drawing. You can also configure the tools the way you like it.

The omissions are few and are more than made up for by, for instance, a decent gradient tool, usable transparency options (ideal for Liquid style cliché art), and even (very) basic 3D shapes. For web designers you can add URLs to objects, export images in SVG format or save art as a self contained web page.

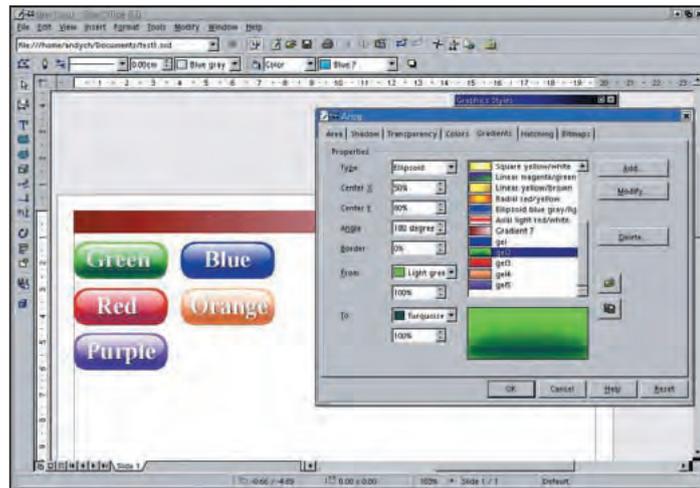
## CONCLUSION

*StarOffice* really is the bees' knees when it comes to business graphics: the native files open seamlessly across platforms, it's rock solid, the tools are comprehensive and the interface elegant. *StarOffice* is not perfect, but you can tell it has benefited from some serious user testing. You may not think you need a drawing app, and if that's the case, just think of this as a bonus. Those of us who need this functionality now have a reasonably priced drawing package with a free office suite. Add *GIMP* and you have every tool you'll ever need. Mostly.

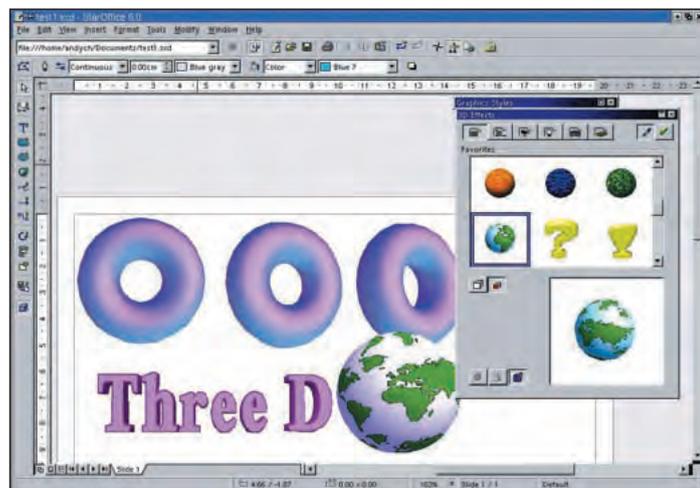
### Star/OpenOffice VERDICT

Has almost all the drawing tools you'll need, from fancy text to DTP features. Not much between the commercial and open options, choose on your need for support and smooth fonts.

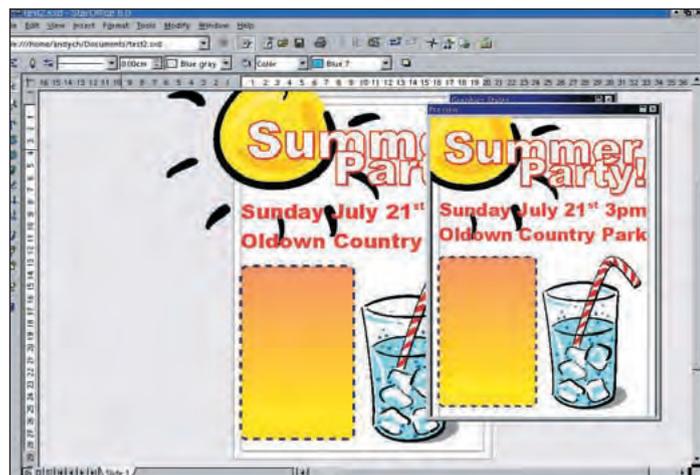
LinuxFormat RATING  
 9/10



*StarOffice* has a good range of preset gradients, or you can set your own!



There are a selection of 3D objects to enliven your documents.



The preview tool is useful for testing bleeds, etc.

## KONTOUR

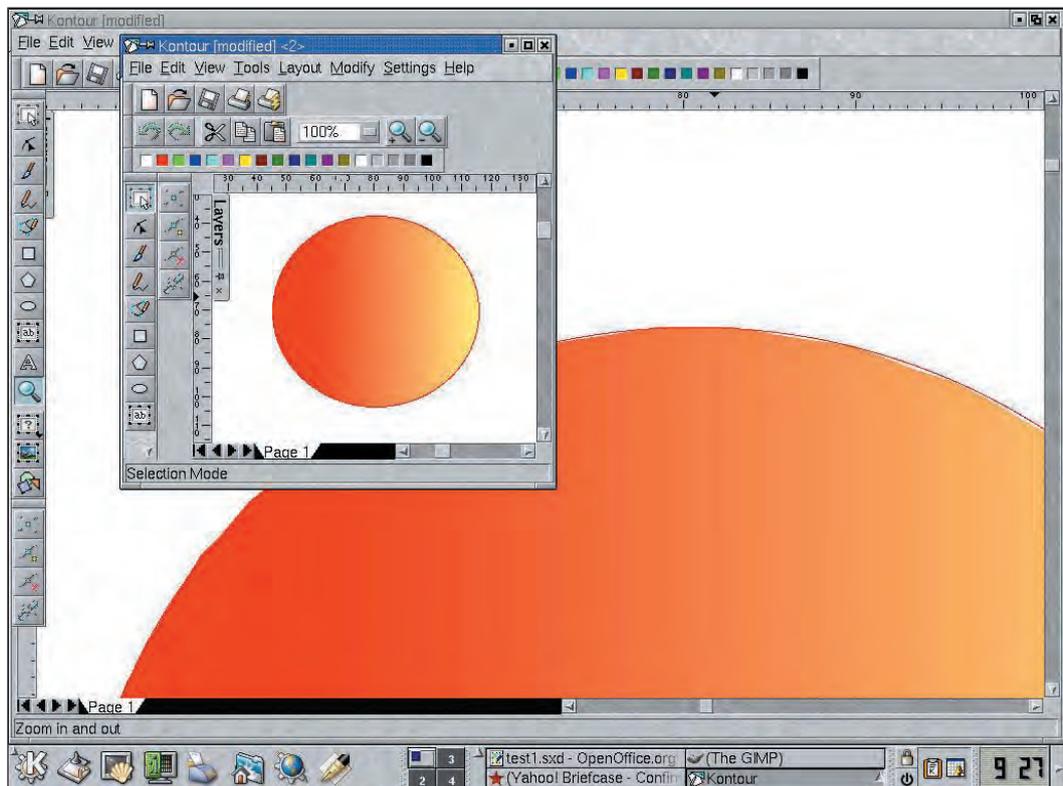
*Kontour*, formerly known as *KIllustrator*, is the vector drawing tool that comes as part of *KOffice* and is the most *CorelDraw*-like of the applications on offer with tools on the left and a range of configurable options along the top of the window. On the tool bar you'll find all the usual primitive shapes, node editing tools and text selectors. Fill and line options, which would make sense here, are instead accessed by right clicking on an object. This is fine, but means you can't preview the effect of a change without first 'applying' it and closing the dialog. It's a bit long winded.

I encountered a similar problem when it came to creating circles or squares: there's no way to constrain the aspect ratio of an object as it is drawn. You can go in to 'Modify' 'Dimensions' to create your perfect circle but, again, it's extra work and lost productivity.

*Kontour* covers all the drawing basics well but I experienced a few problems as drawings increased in complexity; for instance the program throws a segmentation fault if you try to rename a layer and slows to a crawl at stupidly high zoom levels (above, say, 4000 percent) before crashing.

Also on close inspection of the screen image, the drawn curves betray some irregularities. This problem is most apparent if you have a graduated fill on an object, but output, either printed or to file, is not affected. In fact, output is very good, but it would have been nice to see some more options on the file format dialog - no EPS or CDR is a real disappointment.

Integration with the rest of *KOffice* is great, you can insert any other element into your drawing with little fuss and edit it in place. The interface is intuitive, attractive and, for anyone familiar with KDE, instantly



At high magnification levels you can see irregularities in the screen rendering.



You can split the view to get an overview of your closeup work.

recognisable. The 'view' options are also extensive. You can open a new window on an image so you could have an overview and a close up shot for the fiddly work.

This is okay, but each window has a complete toolset and menubar on it

making the screen look cluttered, and also you can only have one window focussed at once.

The better option is to select 'View' 'Split View' which then splits the window in two horizontally or vertically. You can the split your

split view if you like, to add a further level of control, even if each window is an extra drag on performance.

## CONCLUSION

The problems I have encountered using *Kontour* are a shame, because without them this has all the hallmarks of a quality application.

The good news is that *Kontour* is still in very active development so it's definitely one to watch for the future. For now, however, if you're going to attempt anything complicated, stick with [OpenOffice.org/StarOffice](http://OpenOffice.org/StarOffice)

### Kontour VERDICT

Marred by some unfortunate bugs, but still an intuitive, able application. It just needs a bit more development.

LinuxFormat RATING  
6/10

## OpenOffice.org

The big differences between *OpenOffice.org* and Sun's commercial product are in font rendering and the provision of clip art, but pretty much anything *StarOffice 6* can do,

*OpenOffice.org* can do as well. The obvious distinction for corporate users is the availability of technical support that the commercial version offers, but for anyone else I would

recommend *OpenOffice.org* on the graphics front. It doesn't look quite as good on the screen but in printed tests I've found no difference in the output; both have been excellent. If

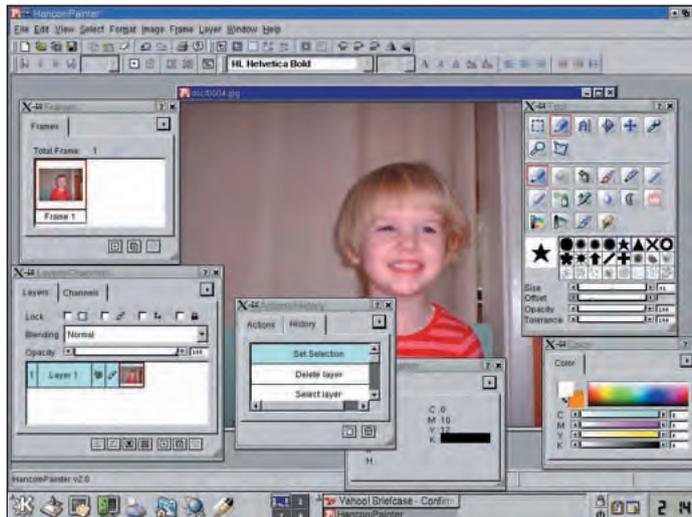
you're a user on a tight budget who needs cross file- and platform-compatibility (or need to support Apple Macs), it's really a no-brainer: download *OpenOffice.org* now.

## HANCOM

The graphical facilities in Hancorn's offering are split into two, the first providing a basic set of drawing tools embedded within the other applications, and the second a fully featured *Photoshop* clone which is more than capable of handling most design tasks. It seems like overkill at first, but as soon as you begin to use *Painter* you'll realise that, as with *The GIMP*, you can hide any redundant options. This means you can simplify the layout to start with and then add elements as you become familiar with the app's depths. The concentration on bitmap editing, however, is mirrored in the lack of any sophisticated vector app; you can add the odd square, circle or text effect, but *Illustrator* style design is definitely out.

*Painter* makes most sense for website graphics, with incredibly useful tools and output facilities on offer. The headline features are Multicrop – which lets you slice up an image into columns and rows and then output each one as a separate file – and jpg export which gives an indication of how long your image will take to download as various compression/quality ratios and connection speeds.

There is also a fairly good set of tools for creating animations including the ability to create frames from any



**Painter is a fully featured Photoshop clone.**

layer, so you can build your animation up in the traditional onion skin way (with transparency set so you can see the previous frame) and then save each layer out as a discrete frame.

The history menu is an undo button on steroids (the undo button, by the way, can be configured from 1 to 100 actions). Every action is recorded allowing you to roll a project back to before it all went horribly wrong. The Actions palette which allows you to record and then play back a series of steps macro style.

*Painter 2.0* offers something that *GIMP* users have been crying out for: CMYK support. *Painter* can generate

or edit images in colour depths from monochrome to 32-bit and all points in between. What's more the batch processing option can save hours of work if you're preparing a set of images for a single document. The application takes a directory full of images and can change format, colour depth, resolution, apply an effect (it can do anything that is recordable as an action) and then save the new versions out into a specified directory or overwrite your originals.

Integration with the other apps in the suite is not so good, for instance you can't launch *Word* from within *Painter* or insert a *Sheet* file into an

illustration, but this is partially compensated by the *HancornShell* which manages to be quite pretty (in a chunky, Tweenies sort of way).

## CONCLUSION

*Painter* would be a worthy purchase on its own. It's capable of everything from basic photo-editing for digital camera fans to sophisticated montages using the filter layers and masks. But where this application really excels is in the field of web graphics. The Multicrop tool is genuinely useful, the export dialog for saving jpegs should ensure your images are optimised and the ability to generate animated gifs is a bonus. The only things missing are automatic rollover and image map generation, but then you can't have everything.

It's not *Photoshop*, but for image editing it's really all you need. But does an office suite need a bitmap package this powerful or would *Hancorn* be better off adding more extensive drawing tools to other apps?

### Hancorn VERDICT

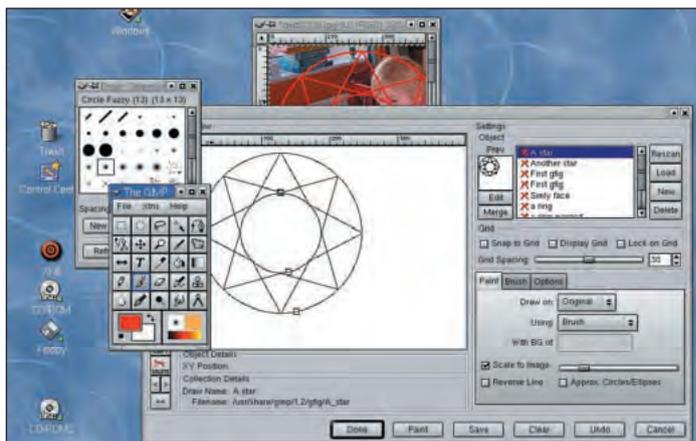
An excellent app that may seem out of place in an office suite. However, for web graphics, it's the dog's.

LinuxFormat RATING  
 8/10

## GNOME OFFICE – THE GIMP

What can you say about *GIMP* that hasn't already been written? It's the Beatles of Linux apps, capable of everything from removing red eye to

creating multilayered montages, it's an image editor that with a little work could seriously put the wind up the likes of Adobe, it's not the most appropriate application for inclusion in an office suite, but what the hell: it should be bundled with everything!



**Gfig is an unusual solution to the 'vector app' problem, but it works.**

*GIMP* has all the tools and features you'd need for day to day image editing including layer support, an excellent gradient fill tool, and a massive range of brushes and textures. The killer feature though is *GIMP*'s extensibility. You can scour the web for plugins, Perl scripts, new brushes and gradients. If you really find you need a vector drawing application, look under filters>render>Gfig to find a basic but useable applet. Sure it's not the best thing in the world, and objects are uneditable once committed to paint (make sure you save your objects separately and 'paint' them on to a new layer to get round that problem), but for basic jobs it's more than capable. *Gfig* can create a range of shapes including spirals, stars, rectangles and ellipses and then render them to your image in a brush and colour of your choice. Rather

than being based on a application-centric Windows paradigm, *GIMP* works around documents and everything is given its own space on your desktop, this can lead to a cluttered taskbar on heavy jobs, but also means each tool is a click away.

With the exception of CMYK conversions, *GIMP* can do almost any purely graphical task you can throw at it. It's not the ideal office suite drawing app, but combine it with *OpenOffice's* strong vector skills and integration and you have a killer package.

### The Gimp VERDICT

*GIMP* is the poster boy for open source applications development. You probably already have it installed.

LinuxFormat RATING  
 9/10

## ANYWARE

Nostalgia is not a bad thing, but is inappropriate in computer apps, so it's a surprise, in an age of enormous hardware and software power to be confronted with a bitmap editor that harks back to the days of the BBC B.

The main work window is laid out in the usual fashion with a toolbox on left and menu bar, file and word processing tools along the top. Just above the scrollbars on the right you'll find a selection of slideshow icons that allow you to control a presentation from within *Graphix*. This is useful for previewing presentations and demonstrates the level of integration achieved in the suite, but like many of the tools on offer they suffer from being obscured by the unintuitive interface – it took a while before I even noticed they were there!

The toolbox features the usual array of primitive shapes, text tool, and fill/line selectors. This feature set is okay but if you've used anything beyond *CorelDraw 3* you'll find yourself hitting the limits fairly quickly. I like the fact that the fill and line tools are so readily available rather than hidden away in flyout or menu but there are so few options, especially when it comes to colour selection, which seems a touch old fashioned.

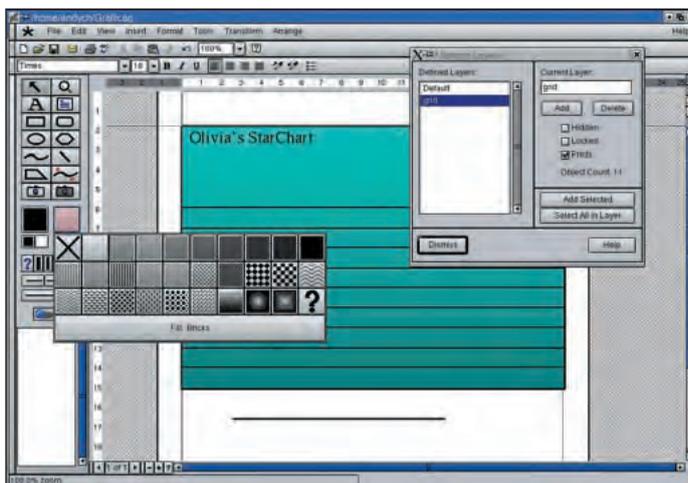
You can select foreground or background colour and pattern (which features 3 basic gradients and lots of halftone style patterns) but there's no complex gradients, no transparency, no multicolour tiled textures. The line options are more extensive, featuring various weights, line style and arrowheads. There's also a nice option to define objects that are inserted into line intersections on other shapes.

One thing that works well is the digitiser options. Here you click either colour or monochrome, and marquis a selection of your image to be rasterised. You can then double click the resulting image to edit it in the awful bitmap editor. This would be useful for creating image slices for web navigation elements but, again, the options – such as removing the thick black border on your selection – are disguised by the interface.

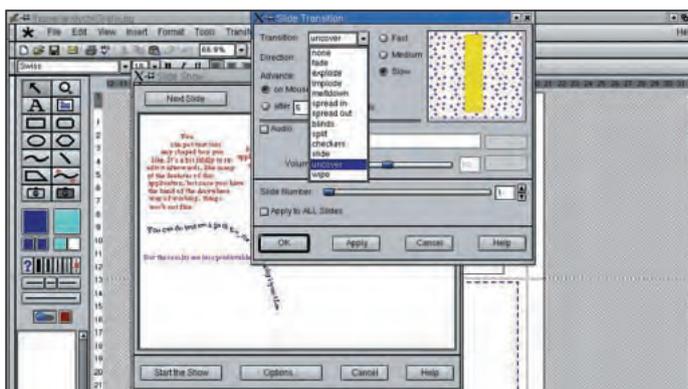
The last section of the toolset is the shadow/emboss button, which does exactly as it says and would be

## Comparisons

Feature	StarOffice/ OpenOffice	KOffice	Hancom	Anyware	GIMP
Vector	yes	yes	no	yes	yes
Bitmap	no	no	yes	yes	yes
Import	bmp, dxf, emf, eps, gif, jpg, met, pbm, pcd, pct, pcx, pgm, png, ppm, psd, ras, sgf, sgv, svm, tga, tif, wmf, xbm, xpm	jpg, kil, png, svg, wmf, x11 bitmap	gif, hif, jpg, bmp, xbm, xpm, png, pbm, pgm, ppm, pcx	im, ag, ap, cgm, dxf, emf, eps, fax, gem, gif, hpgl, ilbm, jpg, mpt, bmp, pbm, pgm, ppm, pcx, pict, pict2, png, ppt, rgb, rs, tga, tif, wmf, wpg, xbm, xpm, xwd	avi, bmp, cel, fits, fli, fax, gbr, gif, gih, gzip, gicon, hrz, jpg, mif, mpeg, pat, pcx, pix, png, pnm, psd, psp, eps, sgi, rs, tga, tif, url, wmf, xbm, xcf, xwd, xpm, bzip2, xjt
Export	bmp, emf, eps, gif, jpg, met, pbm, pct, pgm, png, ppm, ras, svg, svm, tif, wmf, xpm, html	eps, kil, ppm, svg, xpm	if, gif, jpg, bmp, h xbm, xpm, png, pbm, pgm, ppm, pcx	ag, ap, im, cgm, eps, fax, gif, jpg, bmp, ppt, pbm, pgm, png, ppm, rgb, rs, tif, wmf, xbm, xpm, xwd	avi, bmp, cel, fits, fli, fax, gbr, gif, gih, mif, gicon, hrz, jpg, mpeg, pat, pcx, pix, png, pnm, psd, psp, eps, sgi, rs, tga, tif, url, wmf, xbm, xcf, xwd, xpm, bzip2, gzip, xjt



The fill options are somewhat plain.



Integration with the presentation module is especially good,

useful for creating web buttons or adding a little panache to a headline. You can select the distance for the drop shadow and its colour but there are no options for transparency.

Layer support is, again, good to a point: you can hide, lock, make non-

printing (which is good for design notes), but you can't adjust the order of layers, change transparency or delete a layer without first deleting its contents.

First impressions of *Graphix* aren't good. It looks a bit clunky and the

more useable options are hidden away and not duplicated in the menu options. To make a perfect circle, for instance, would be simple if you could constrain the aspect ratio with the **Ctrl** key, but you can't; you have to double click the object and edit its dimensions numerically. This is okay for accuracy, but doesn't make for a fluid design process (it's a similar story if you want to rotate an object). It may be that the developers were hoping to introduce users to the system gently, but it may lead to some trying out the demo and deciding the app is too limiting.

For bitmap editing you'll be better off using practically anything else, so dedicated photo tweakers are best off avoiding this, but if you do occasional, undemanding illustration work *Graphix* should provide the tools you need, albeit in an interface that requires a bit of extra work.

## CONCLUSION

When it comes down to it, *Graphix* is adequate. But to be a contender against an app as good as *StarOffice* it needs some serious development.

### Anyware VERDICT

The ugly interface hides a useable application with good integration. Just don't try to edit a photo.

LinuxFormat RATING  

 6/10

# Presentation Tools

If your job involves standing in front of a group and talking, good presentation software will help you illustrate whatever points you need to make. **Neil Lucock** reviews the tools supplied in the office suites.

When a company is considering a move to Linux there are many factors to be considered. I have three years' worth of Powerpoint drawings on my hard drive at work. Whatever suite you choose, it ought to be able to import your old artwork.

For each presentation package, I attempted to open an *Office 2000 Powerpoint* presentation that I made for Railtrack NW Signaller's Safety Briefing and a presentation that came from Railtrack HQ that has animated transitions and sound effects.

The cost benefits of a move to Open Source can easily be wiped out by high training costs. If all you need is a brief orientation or can get away with training just a few staff to act as mentors to the rest, you will benefit from the move to Linux. "Ease of use" and "Unix" have not always fitted into the same sentence well. I'd want software that everyone can understand and use right away, in other words similar to what they are used to: *Powerpoint*.

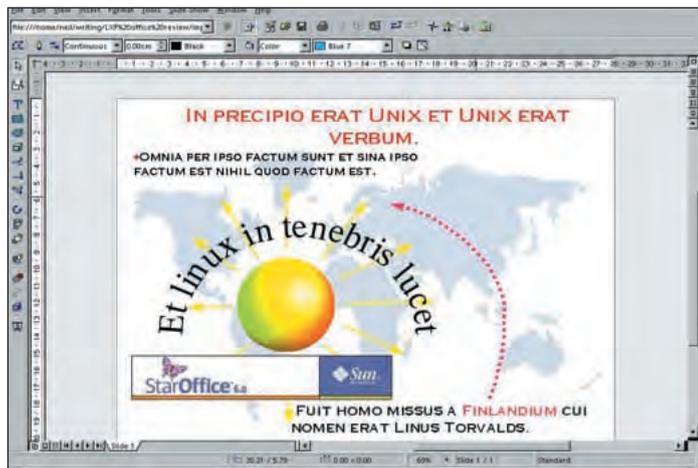
I'll also briefly discuss the features each presentation package offers, anything that makes it stand out from the rest and any other qualities that impressed me.

The presentation suites were reviewed on a 700MHz PIII with 512MB of RAM running KDE 2 in Mandrake Linux 8.2

## STAR OFFICE 6 IMPRESS

*Star Office* starts with the word processor. You have to then use the File>New menu to get to *Impress*. The wizard helps you to create the basics. You can use one of the supplied templates, open an existing presentation (with a preview window to show what you are getting) or start with a blank one. You get a nice selection of templates (in effect, suggested outlines for different kinds of presentation. You just put the details in).

*StarOffice* has been around for a few years and is available for Windows, Solaris and Linux. Sun has obviously listened to their user's feedback, what you have is a polished product that is easy to use. If you can already make presentations, *StarOffice* is going to be no problem at all. I instantly felt at home. It comes with a lot of reasonably good clip art, some photographs and a gallery browser so you can see what's available. It supports layers, so you can act on just the items on the layer.



There's a 3D option, allowing you to colour and illuminate a variety of basic shapes. The range of graphic import options is good, it even loaded DXFs made in a 3D program. Their Fontwork is nice to have. It allows text to bend along a curve. You can draw curves, all manner of lines, arrows and connections. It will make simple webpages for you. You can make hyperlinks. The list of features is huge.

Was anything wrong? Is it a perfect piece of software? The Help system says that it allows **Ctrl-C/Ctrl-V** from other parts of Linux. I tried to paste some text opened in a text editor, but it would not do it. If you open the text in their word processor, it does work. I would have liked the Modify>Arrange>"Bring to front" (and the other options on that menu) from the drawing part of the program to have been included in the presentation menus. However, the keyboard shortcut works fine.

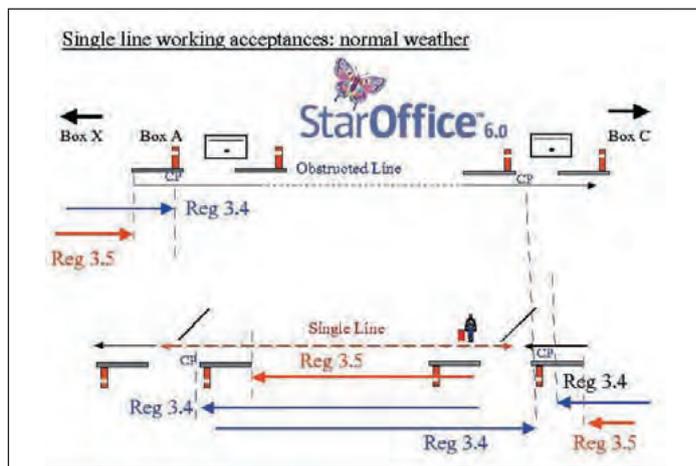
How did it import the *Powerpoint*

**StarOffice Impress makes it easy to produce attractive work. The 3D and text effects are not found in the other presentation programs. If you want it to look wonderful, Impress is the one to use.**



**Cyrillic, Hebrew and Arabic characters are no problem for StarOffice Impress**

files? *Star Office* rendered the test file perfectly. It also loaded and ran the HQ presentation with no problems. If you were in the audience, you would not have known that someone was using Linux. *StarOffice Impress* is a great piece of software. They charge \$75.95, about £50 to £60 in real money. Excellent value, in my opinion, for something that I think is superior to *Powerpoint*.



**Impress lives up to its name, a perfect display of the *Powerpoint* test file. I was impressed. On the animated parts of the other test file it displayed it better than Windows NT 4 does!**

## Star Office 6 VERDICT

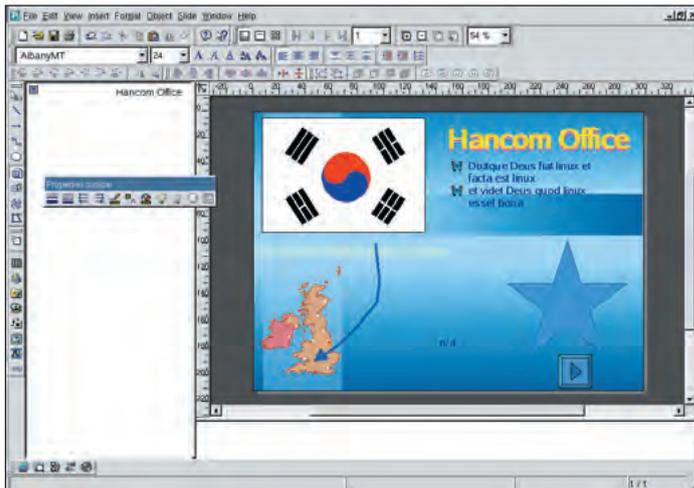
Beautiful output, easy to use and feature-rich – when they named their software *Impress*, they got the name exactly right.

LinuxFormat **RATING**  

 A rating bar consisting of 10 red slanted rectangles, with 9 of them filled, representing a 9/10 rating.

## GNOME Office

The GNOME web site gives details of *Achtung*, a presentation utility, but it is in the early stages and development seems to have stopped according to one of the developers. GNOME are also contributing to *OpenOffice.org* (which is based on *StarOffice* code).



**Hancom Presenter** – clear, simple and a joy to use.

## HANCOM OFFICE 2 PRESENTER

The review software was an evaluation copy with no Help files or templates, but it impressed me from the start. It even put a *Hancom Office* entry into the Kicker menu in KDE. It installs a "Hancom Shell" that you can use to quickly launch whatever part of the suite you need. It does not look quite as polished as *StarOffice*, but once you



**Hancom's Shell for launching different parts of the suite.**

explore, you realise that *Presenter* is high quality, loaded with a lot of good features. Its wonderfully easy to use. It asks you for your slide layout choice, any background or template you want, and opens your new presentation.

I noticed a few jagged edges on the background graphics where there were curves. The menus are context-sensitive, in that the parts you cannot use are greyed-out. The toolbars are all detachable (*StarOffice* tools are fixed) and you double right click to put them back. The copy and paste commands worked from *Hancom Word* but not from *Linux Advanced Editor*. You get a lot of good quality clip art, you can ungroup it and alter it easily. (The evaluation copy would not let me view it until I became Root.)

The whole program is well put together with nice touches. *E.g.*, the slide show rehearsal has analogue and digital displays on the clock to show you how long you have taken. When you have finished, it displays a box telling you how long you took.

*StarOffice* displays just a small digital display with no time report after you finish your rehearsal. *Presenter* is so packed with features, the review is in danger of becoming a long list of what it can do. Insert a .wav, .mp3 or .mpeg? Not a problem. Want to use a butterfly as a bullet point? Rotate an object 90 degrees? Click the icon.

Just occasionally you get a cryptic error message. There's a nice "Edit" button, for when you want to alter the curved line you have drawn. Try to use Edit with a star polygon shape and it tells you to convert it to a curve first, but does not tell you how. I tried to paste it into the Painting part of the program, but it did not want to know.

And the *Powerpoint* import? You would have to spend a few minutes editing the artwork before you could use it, but it's no worse than between different versions of Windows. I was pleased with the HQ presentation, which also imported well.

*Hancom Office*, at around \$50, is excellent value. Why are people trying to make MS *Office* run under Linux when there are apps of this quality around? Far better than *Powerpoint*.

### Hancom Office 2 VERDICT

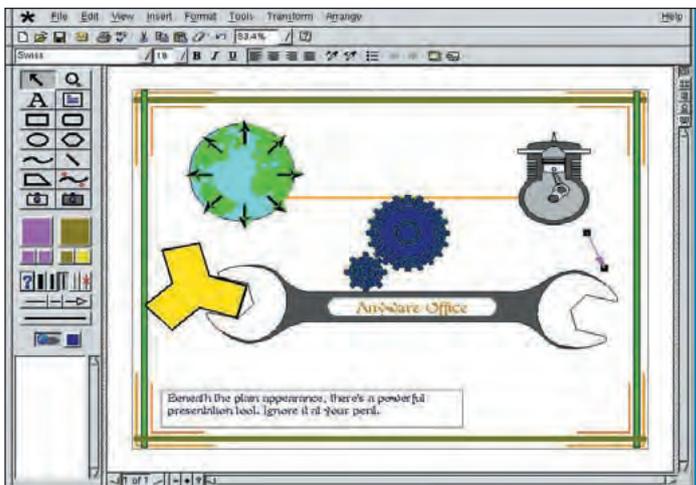
Absolutely brilliant, packed full of excellent tools that make work a pleasure, intuitive and well designed.

LinuxFormat RATING  
9/10

## ANYWARE OFFICE PRESENTS

I was initially disappointed at the dated interface. It looks and feels ten years old. Whereas the other suites in the review have benefited from long development times, *Anyware Office* appears to have never changed. Still, appearances can be deceptive.

There's obviously a lot to the suite. When started, *Presents* gives you various options, just like the other suites. Template presentations are provided, you have a choice of different backgrounds (there are 17), everything you need. Once you start using *Presents*, you realise that all the tools you would expect are there. There are 99 fancy borders.



**Anyware Presents** has loads of clip art, you can put together a reasonable presentation without being able to draw. Note the dated interface.



**Anyware Office's icon bar allows launching of different parts of the suite.**

I found *Presents* awkward to draw with. In particular, it was sometimes difficult to select things. I put a piece of clip art in that was too big. There's a handy keyboard shortcut to shrink it, but, it shrinks a copy, the original one is still there! There are great tools like "Copy Attribute" and, if you take the time to learn *Presents*, you can produce good presentations. The more I used it the more I realised that it is better than it appears to be. The tools are there, they just not obvious.

Take the available colours as an example. It looks like there are just 17 colours to make lines etc. Click on the right part of the program and you'll access all the other colours. It's not poorly featured, it's just different from what you might expect. There's a huge range of import options. Just name a format to open and it will do it.

*Presents* displayed the *Powerpoint* test file perfectly, although it was a bit

slow loading. I tried the presentation from HQ. It crashed and would not start again until I closed KDE and restarted the window manager.

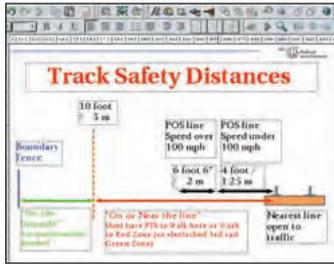
Spend some time with *Presents* and you'll see it is a workman-like utility that gets the job done. There's nothing wrong with it; *Star* and *Hancom Offices* just do the same tasks (and more), do them easier and are both cheaper. You'll need to take some time to learn *Anyware Presents*. If the rest of *Anyware Office* has tools that you really must have, then *Presents* will be adequate.

### Anyware Office VERDICT

It's a competent piece of software but you'll need to learn it and it will never produce work as beautiful as *Impress*.

LinuxFormat RATING  
6/10

# UltimateOfficePresentationTools



**KPresenter did not load the Powerpoint test file, but is capable of good results, although here adding line boxes around the text.**

## KPRESENTER (VERSION 1.1.1)

KPresenter does not have the huge range of clip art, backgrounds or tools that the others have, but it is very usable and you can get acceptable results easily. If you want to quickly make a few slides that don't need flashy effects, it's fine. My copy failed to import the Powerpoint files, so that's something that still needs work.

### KPresenter VERDICT

Included in KOffice, so you'll already have it. Easy for everyday use but lacks the import abilities and features.

LinuxFormat **RATING**  
 7/10

## OPEN OFFICE.ORG IMPRESS

StarOffice 6 (SO6) is the latest version of the suite Sun bought from StarDivision in 1999. OpenOffice.org (OOo) suite is really SO6 code released under GPL and the Lesser GPL. Sun can only release code they own and SO6 contains elements that are licensed from other companies. So, in OOo some features are missing, but they fall into the nice-to-have category rather than being essential.

If you want all the time-saving backgrounds, fonts and slide designs, SO6 has them. The templates (offered when Impress starts) are missing, the slide designs are missing although the slide transitions and Autopilot are there. Autopilot just makes the presentation with a plain background. In SO6 it will give a choice of twenty templates to use for Brainstorming, Company Reports and other useful business presentations.

## Comparisons

Feature	Star Impress	Open Office Impress	Hancom Presenter	Anyware Presents	Kpresenter
<b>Text</b>	Justify Line spacing Bullets Font preview Arabic/Hebrew	Justify Line spacing Bullets Font preview Greek/Cyrillic	Justify No Bullets Font preview Everything	Justify No Bullets No Latin & Greek	Justify No Bullets Not all Latin/limited
<b>Basic shapes</b>	Yes	Yes	Yes	Yes	Yes
<b>Bezier curves</b>	Yes	Yes	Yes	No	No
<b>3D shapes</b>	Yes	Yes	No	No	No
<b>Lines/arrows</b>	Yes	Yes	Yes	Yes	Yes
<b>Connectors</b>	Yes	Yes	Yes	No	Yes
<b>Sounds</b>	wav	wav	"wav, mp3"	Au	No
<b>multimedia</b>	mpeg	No	mpeg	No	No
<b>clip art included</b>	Yes-lots	Yes-some	Yes	Yes	No
<b>Bookmarks</b>	Yes	No	No	No	No
<b>Header/footer</b>	Yes	No	Yes	Yes	Yes
<b>Rehearsal timer</b>	Yes	Yes	Yes	No	No
<b>Spellchecking</b>	Yes	No	No	Yes	No
<b>Thesaurus</b>	Yes	No	No	Yes	No
<b>Footnotes</b>	Yes	No	Yes	No	No
<b>Hyperlinks</b>	Yes	Yes	Yes	No	No
<b>Webpage</b>	Yes	Yes	Yes	Yes	Yes
<b>Find/replace</b>	Yes	Yes	Yes	Yes	Find
<b>Slide sorter</b>	Yes	Yes	Yes	Yes	Yes
<b>Object animation</b>	Yes	Yes	Yes	Yes	Yes
<b>Layers</b>	Yes	Yes	No	Yes	No
<b>Macro</b>	Yes	Yes	*	Yes	No
<b>Notes</b>	Yes	Yes	Yes	Yes	No
<b>Permissions for exports</b>	No	No	No	Yes	No

Notes "Star includes Arabic & Hebrew (right to left), Greek & Russian scripts. "  
 Hancom Office includes nearly every font on the planet.

\*" The review copy from Hancom had no help files. I was unable to find any macros, but they could be there."



**OOo Impress has the same tools as SO6. If you want to use the Fontwork more than once in a slide, convert the parts you are happy with to a bitmap.**

When you open OOo you get the same interface as Impress in SO6 gives you. (It opens presentations made in SO6, it uses the same XML file format). The foreign fonts have been removed, but do you really need Japanese fonts in your presentation? What you get is the terrific tools included in SO6 but none of the work done for you. Want a background to your slides? Make one. The orange

gradient background (in the screenshot) took about five mouse clicks, the 3D effects and font effects are all there. It's just as easy to make something attractive as it is in SO6. The help files are there, it printed correctly, but the spellchecker lied. There are other features missing – see the comparison table for details.

Both Powerpoint test files imported perfectly, exactly the same as they did in SO6. This fact alone makes OOo a major resource for anyone trying to advocate a move to Open Source software. If you can manage without the range of useful extras provided by SO6, then OOo is an excellent choice.

### OpenOffice.org VERDICT

It lacks a few things that the others have as standard, but makes up for it by being the best value for money.

LinuxFormat **RATING**  
 9/10

## Conclusion

I have a hard time choosing between StarOffice and Hancom Office. StarOffice makes really beautiful presentations, but Hancom Presenter is both an excellent piece of software and a little cheaper. It did not import my test files as well as StarOffice did, but the rest of the program is a joy to work with. If possible, try each suite before you buy to see which suits your style of working best.

I found Anyware Presents to have great depth, but it was let down by a dated interface that was not intuitive. There was neither the pure ease of use that Hancom Office had, nor the beautiful output from StarOffice.

KPresenter is easy to use and surprisingly effective, but not as well developed. OpenOffice.org's Impress has most of StarOffice 6's qualities, including Powerpoint imports but none of the cost. If you can manage without the parts that have been removed, you still have a powerful presentation tool. The choice is yours.

# The verdict

The votes are in and the foreman is ready – it's time for the verdict on Linux Office Software.

We have covered all the main components of the individual suites, but of course, there are other things to be taken into consideration. Documentation for one thing is important – clear instructions and online help across the whole range of applications is a must for ease of office deployment.

There is also the question of reliability (you don't really want your presentation to crash in the middle of an important meeting do you), not to mention how well integrated the software is. We have given scores on these factors in the table on this page.

If you have read the preceding pages then it will come as no surprise to see that StarOffice has claimed the top gong. As well as being outstanding in its ability to read MS Office documents, each element of the suite is also feature packed and easy to use. I must say that we ourselves were surprised at how good it has become since the OpenOffice.org project was set up. What used to be a slow, clunky suite that didn't sit happily on the Linux OS has turned into a real star performer.

Both the Anyware and Hancom solutions were good too, and it may be that, depending on your needs, these provide a better solution. They have their strengths and weaknesses – HanCom Presenter was a revelation, while Anyware's word processor was top class.

## First place Star Performer



With an impressive final score, and top marks in all categories, the runaway winner is StarOffice. It gets the nod over OpenOffice.org mainly because of the extra features, support and documentation that couldn't be provided by its free software counterpart.

Earlier versions of StarOffice had their problems, and while it isn't by any means perfect, it is hugely improved – the decision to open source it has obviously paid off big time – it is certainly much faster and feels like more of a Linux application than it used to. The consensus seems to be that Sun are charging a great value price for the software too, and the licencing terms are generous.

## Second place OpenOffice.org

With pretty much an identical feature set to StarOffice, the only things missing from its open source sibling are documentation, support, examples and a few extra features. As the software is essentially free, this makes it incredibly good value. While corporates may want the security of a support system (would you bet your business on it?), for the home user OpenOffice.org simply can't be beaten.



KOffice's main benefit is that it is well integrated into the desktop environment, but most of the tools themselves are fairly basic when compared to the others on test here.

A word has to be said about the GNOME tools too. Both Gnumeric and GIMP are excellent pieces of individual software, and may well serve your needs for specific applications. Where they excel, apart from the application features, is the huge range of support

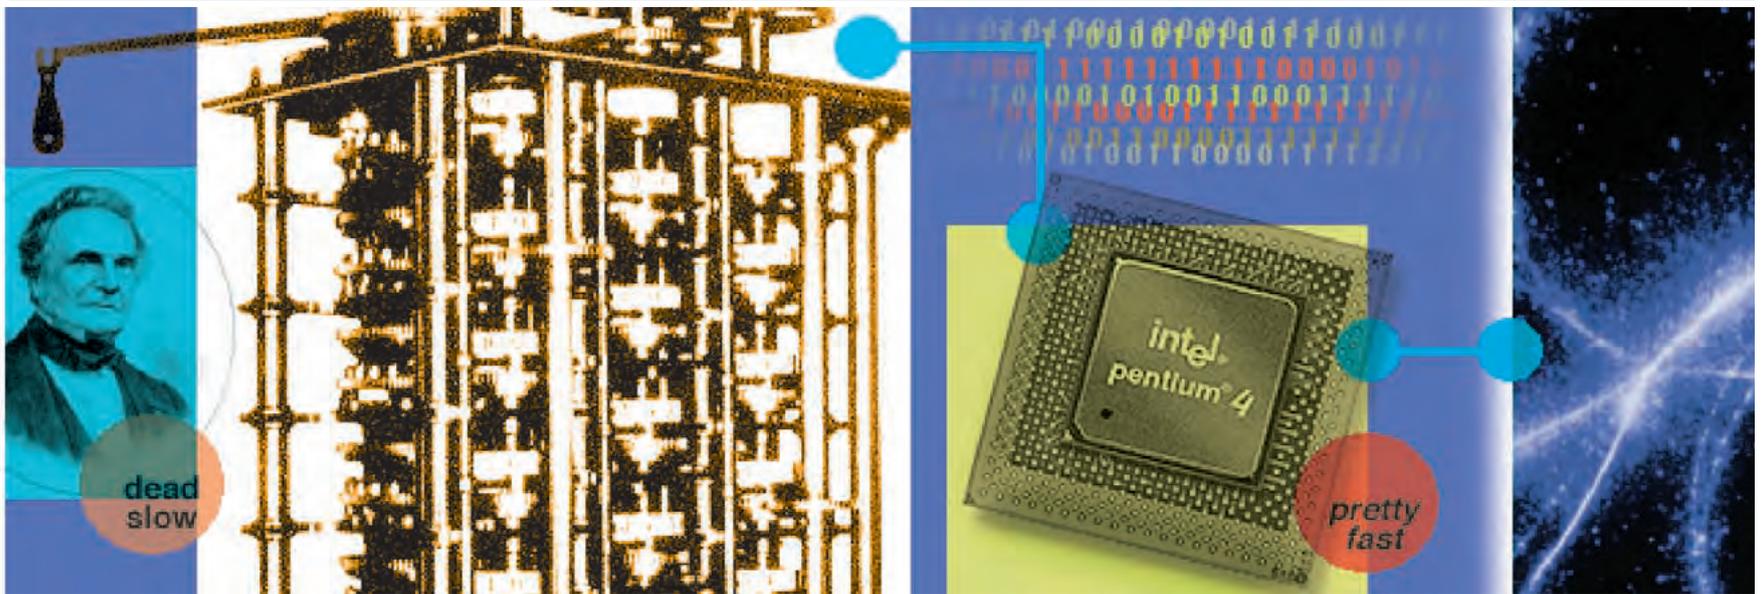
for other document types both of these possess.

Huge though this feature has been, and great though the number of hours of work that have gone into it, there is always more to say. The trouble is that, although office applications are the most commonly used, there are many ways to use them. The beauty of Linux though is that, if you don't believe our findings, you can try out the software yourself.

There are trial versions of Hancom Office and Anyware Office available for download, OpenOffice.org (on LXF29 coverdiscs) is so similar to Star Office there isn't much practical difference, and the KOffice and GNOME tools form part of pretty much every modern distribution. If you have any opinions or experiences you want to share on Linux office software, write to our Mailserver section at [lxletters@futurenet.co.uk](mailto:lxletters@futurenet.co.uk) [LXF](#)

## Overall comparisons

Feature	KOffice	GNOME Office	Star Office 6	OpenOffice.org	Anyware Office	HancomOffice
Version	1.1	various	6	1	2.2	2.01
Web	<a href="http://www.koffice.org">www.koffice.org</a>	<a href="http://www.gnome.org">www.gnome.org</a>	<a href="http://www.sun.com">www.sun.com</a>	<a href="http://www.openoffice.org">www.openoffice.org</a>	<a href="http://www.vistasource">www.vistasource</a>	<a href="http://en.hancom.com">en.hancom.com</a>
Price:	Free	Free	£52.99	Free	\$99(£70)	£49.95
Word Processors	5	7	10	10	9	7
Spreadsheets	5	8	9	9	7	8
Graphics	6	9	9	9	6	8
Presentation	7	-	9	9	6	9
Integration	7	2	8	8	8	5
Documentation	1	varies	9	7	7	7
Ease of Use	6	8	9	9	8	9
Features	5	7	10	10	9	8
Overall	5	6	9	9	8	8



# What on Earth is... QUANTUM COM

**Andy Channelle** puts on his thinking cap to explain computing and dead cats.

## » Oh no, pass the paracetamol I feel a headache coming on!

Don't worry, we'll try to keep it simple, but bear in mind that Niels Bohr noted that if you're not confused by quantum physics, you don't understand it. In fact, some scientists go so far as suggesting that quantum physics cannot be understood, and one must 'simply get used to it'.

## » Yeah, yeah, we all know it's tricky stuff, but what's it got to do with computers and do I need to read about Dirac, Heisenberg and Schrodinger's feline abuse to understand it?

Right, stripped down to the basics, a quantum computer harnesses the physical phenomena unique to quantum mechanics in order to speed up computation far beyond anything possible (now or ever) with a machine bound by classical physics. Quantum computers make use of the properties known as Quantisation, Interference and Entanglement. Quantisation is the bedrock of both

classical and quantum computing and just means that observable quantities come in discrete packages – what Max Planck termed quanta – that vary in steps (like a digital audio signal sampled 44,000 times per second) rather than continuously (like an analogue sine wave).

Interference is a little more complex and is best demonstrated by the double split experiment that Thomas Young devised in 1801 to contradict the accepted Newtonian idea that light was 'corpuscular' in nature. Young devised a simple experiment that suggested light travelled in waves. What he did was shine a monochromatic light at a viewing screen, but disrupted the light's passage with a second screen that had two thin slits in it. This creates the typical striped diffraction pattern caused by areas of 'constructive interference' where the two curved light waves interact – the bright strips result from waves arriving crest to crest, while the dark areas resulted from them arriving trough to crest. Seems simple enough, but it gets more confusing.

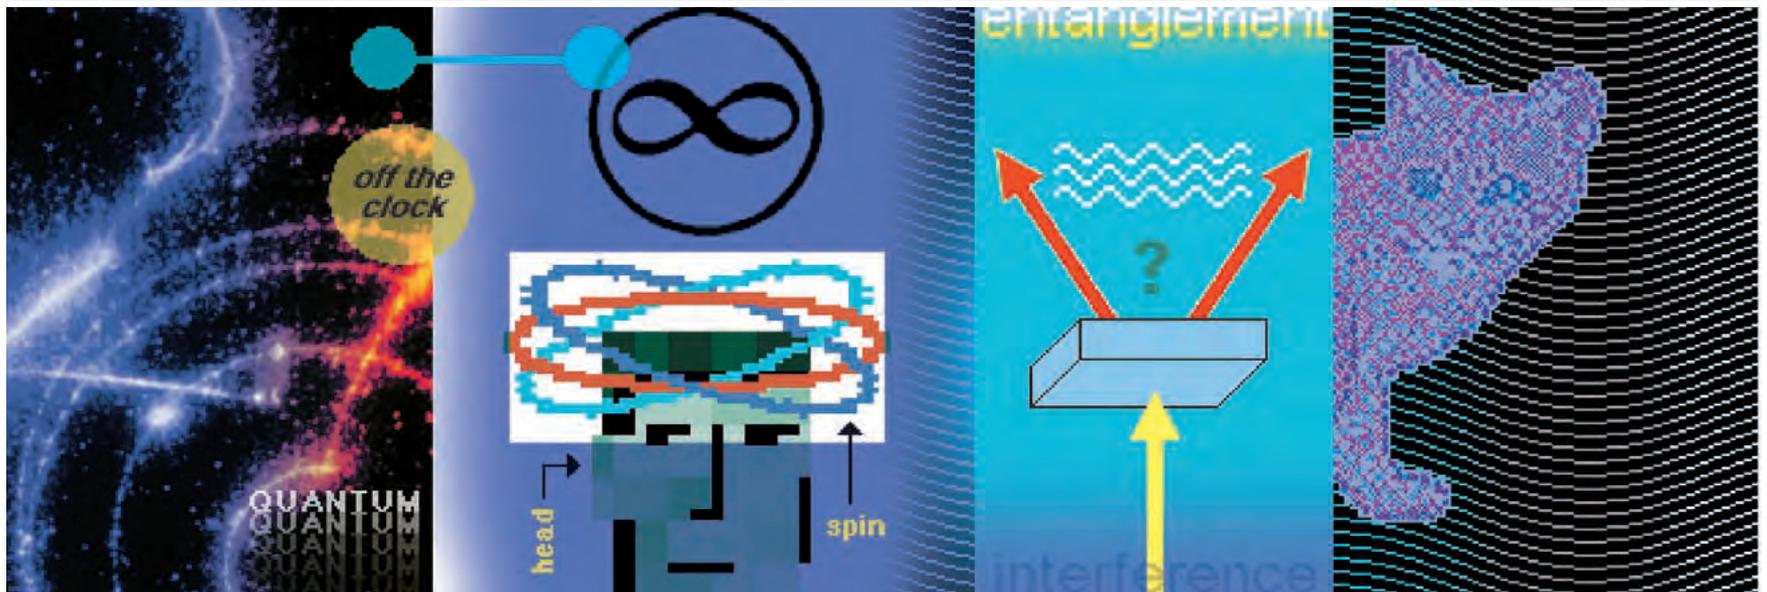
When the technology permitted it, physicists

refined the double split experiment by firing light particles (photons) at the screen one at a time and surprisingly, the same interference pattern appeared, suggesting that an individual photon travelled like a wave through both slits simultaneously. This is particle/wave duality; light behaves like particles and waves simultaneously.

With regards to the second part of your question, do you need to know how a cotton loom works in order to enjoy fresh sheets on your bed?

## » How does a quantum computer differ from a classical computer?

Your traditional desktop PC, despite its speed and size, is built on the same principles as Babbage's pioneering 1832 Analytical Engine and Turing's universal computer: a quandary was punched in and, after doing its steam-driven work, the machine spat out a result. Babbage's machine had 50,000 components, could 'store' 1000 numbers of up to 50 decimal digits long and was the size of a house. Your average Pentium 4 packs about 42 million transistors in something the size of a chocolate bar, but despite these refinements, input goes in one end, electricity travels through a series of switches, and output comes out of the other. So basically,



# PUTTING?

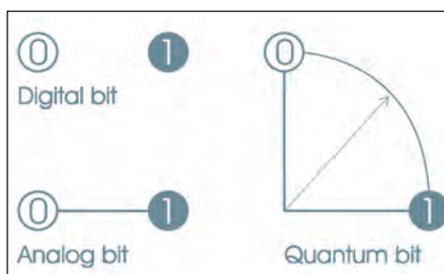
given enough time, memory and output options, any classical computer can do the work of any other. That goes out of the window in the quantum world.

The quantum computer was first seriously proposed in a 1978 paper (published in 1985) by David Deutsch. In that same year Deutsch also published the theory that leads to quantum parallelism – the property that makes these things so fantastically powerful – suggesting that if a single electron can move along a number different paths simultaneously (and it can), a computer working along quantum lines could calculate along many difference routes, thus speeding up the process by an enormous factor.

This nascent technology already has a pair of killer applications: Shor's algorithm (for factorising numbers) and Grover's algorithm (for combing through unsorted databases). What has excited David Deutsch about Peter Shor's work is not its mathematical elegance but the apparent proof it provides of the controversial 'many worlds' theory. Setting out the idea in his book *The Fabric of Reality*, Deutsch said a Shor computer would be able to compute along  $10^{500}$  simultaneous pathways, and as there are only about  $10^{80}$  atoms in the entire visible universe, "physical reality would not even remotely

contain the resources required to factorise it!" So where does this process take place? Deutsch says this happens in other parts of the multiverse. "A quantum computer would be capable of distributing components of a complex task among vast numbers of parallel universes and then sharing the result"

At the heart of any computer has to be information. Classically this is represented by binary digits (bits) which, in a digital system, can be in one of two positions; an analogue bit can exist at any point between the two positions. A qubit exists as a superposition; it is both 0 and 1 at the same time – this is shown as a vector at the mid-point between the two states in the above diagram. In practice this means that a two qubit computer can represent four



Digital, analogue and quantum bits.

numbers simultaneously, while its traditional counterpart would only be able to handle one. Moreover, adding qubits to your computer increases its power exponentially – a 333 qubit computer would be able to use a number greater than all the atoms in the known universe.

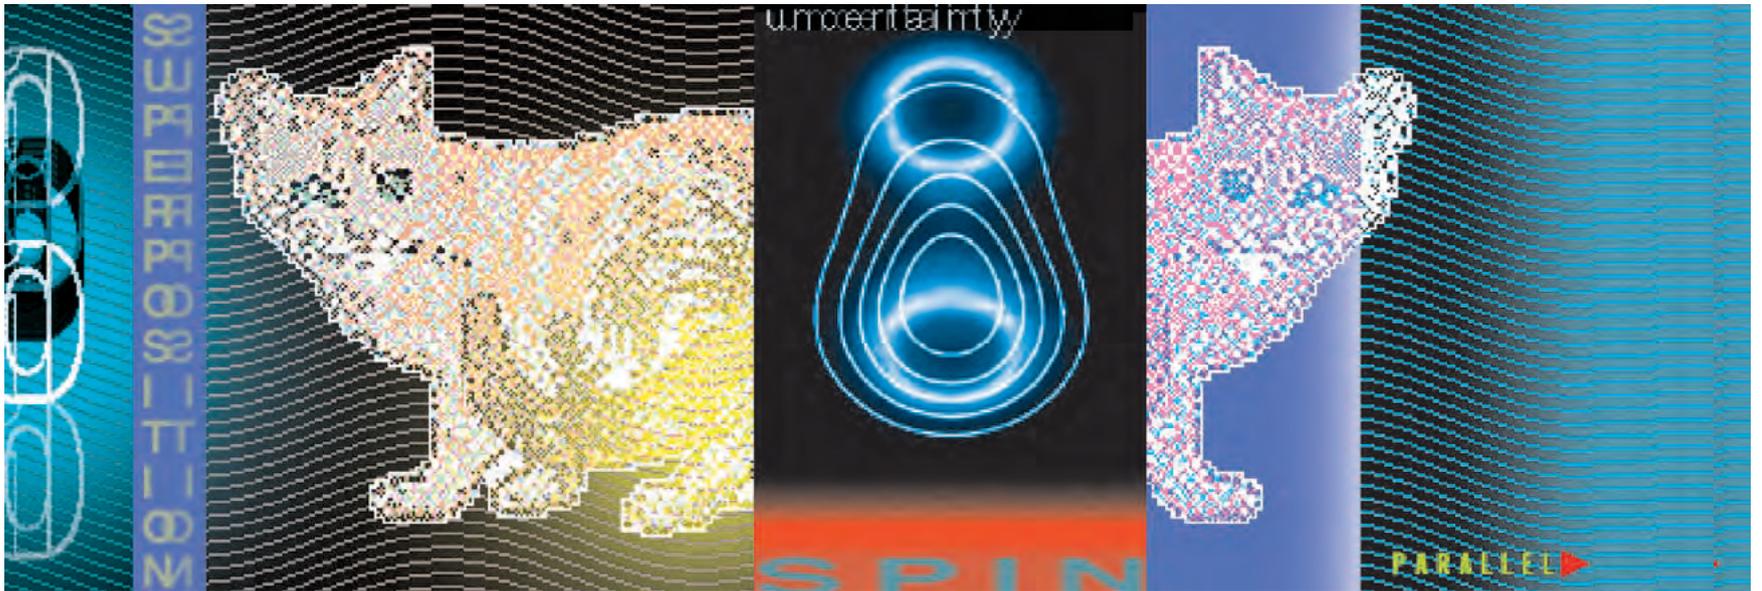
## »» This all sound marvellous in theory, but that's all it is right? A theory?

It's probably not stretching the point to say there has been some scepticism about the whole idea of quantum computing; some pundits have gone so far as to suggest that it would be impossible to implement something as radical as Peter Shor's algorithm in the real world. That hasn't stopped a number of companies (and probably governments, of which more later) throwing enormous sums of cash at the problem though, and in December of last year IBM announced a breakthrough. Scientists at Big Blue's Almaden Research Centre in California had succeeded in fashioning a limited quantum computer capable of factoring the number 15. By today's standards this computer was piffling – and it looked somewhat like a test-tube full of urine – but this was nonetheless a giant step in the right direction.

IBM's quantum computer was created by Issac Chuang, one of the world's leading experimentalists in the field. Chuang and a team of chemists designed and made a new molecule that has seven nuclear spins – the nuclei of five fluorine and two carbon atoms – which are able to interact with each other as qubits. These qubits can then be programmed by radio frequency pulses and detected by nuclear magnetic resonance (NMR) instruments. The successful computer contained approximately a billion, billion of these molecules and required an unprecedented level of control over the spins during the calculation.



# WhatOnEarthQuantumComputing



« The big problem with IBM's system is it is difficult to create molecules with more than seven qubits so the search is on to find a technology that will 'scale' more readily to fulfil the incredible promise of quantum computers. Most promising is the research into quantum dots, which utilise electron spins confined in 'semiconductor nanostructures.'

## »» What's the point of all this power? I still find an old 486 can take care of most of my day-to-day computational needs.

Think of something fast – not just merely nearly fast, but really most sincerely fast. A quantum computer will be faster than that. Julian Brown, author of *Minds, Machines and the Multiverse*, says these machines will make today's supercomputers seem puny: "It would solve problems that will never be cracked by any conceivable non-quantum successors of current computers." But quantum computers aren't going to be for playing games or predicting the weather, in fact finding applications that might make use of the technology is exercising the brains of those in the field as much as the mechanics of the machine itself.

If bets were being taken on the most useful quantum computing apps, LXF would have its cash placed on the square marked Factorisation, the very job IBM's urine-esque creation has made a reality. It was designed to calculate the factors of 15 (which are three and five) which, while not exactly difficult, is a start. Now, take the number 267,023; working out the factors for this number is harder, and the difficulty is increased because both numbers are primes, in this case 257 and 1039. Given a little time your average desktop PC could find this out, but if the number was, for instance, 1,658,805,411,577 you're looking at quite a long job (the factors are 1,278,787 and 1,297,171, by the way). And why is this important? Because enormous numbers and their prime factors is what public key encryption systems such as RSA (the core of PGP) are based on – your public key is the huge number (ideally a couple of hundred digits), your private key its factors. There are

traditional ways of factorising, but they'll still take a long time to break your average public key system.

This sort of calculation will be routine for a program capable of implementing Shor's groundbreaking algorithm, which until the advent of IBM's 'machine' was a theory looking for a practical application. The novelty of Shor's work is that it is capable of factoring numbers in a reasonable time regardless of their complexity, so it doesn't take a genius to figure out who's going to be spending the most money on quantum computation. Imagine, the secret emails of your enemies, competitors and critics laid bare.

## »» Assuming that governments are always a couple of steps ahead, does this mean my PGP encrypted email can already be read by James Bond and his spooky mates?

That's a difficult question to answer. Long after Ronald L. Rivest, Adi Shamir and Leonard M. Adleman (the R, S and A in RSA) had published a paper on their encryption system and received the glory that goes with such a discovery, it emerged that cryptographers toiling away in Britain's GCHQ had made the same leap years earlier. They had to stay in the dark, however, as the British government didn't want everyone else to know they had a secure communication method. So the answer is: probably. Which is fitting for a thing about quantum physics.

## »» So that's the end of secrecy then?

Not necessarily. While quantum computing may be able to crack (in the future) any level of RSA style encryption in a short time, quantum physics also introduces an encryption method which is genuinely unbreakable: the panacea of cryptography and a nightmare for cryptanalysts.

A quantum cryptographic system is based on the public key idea which, according to Todd Larchuck, expert on experimental quantum optics, "attempts to link the security of the system to the correctness of the uncertainty principle of quantum mechanics." Using Heisenberg's uncertainty principle, the system

works on the basis that any information extracted from an encrypted message will irreparably alter or disrupt the system. Therefore Alice and Bob (the traditional sender and receiver of encrypted messages) will be alerted if Eve, the nefarious third side of the cryptography triangle, manages to intercept their message.

## »» Sounds tricky! Is there a catch or can I start sending out quantum encrypted love notes from Kmail?

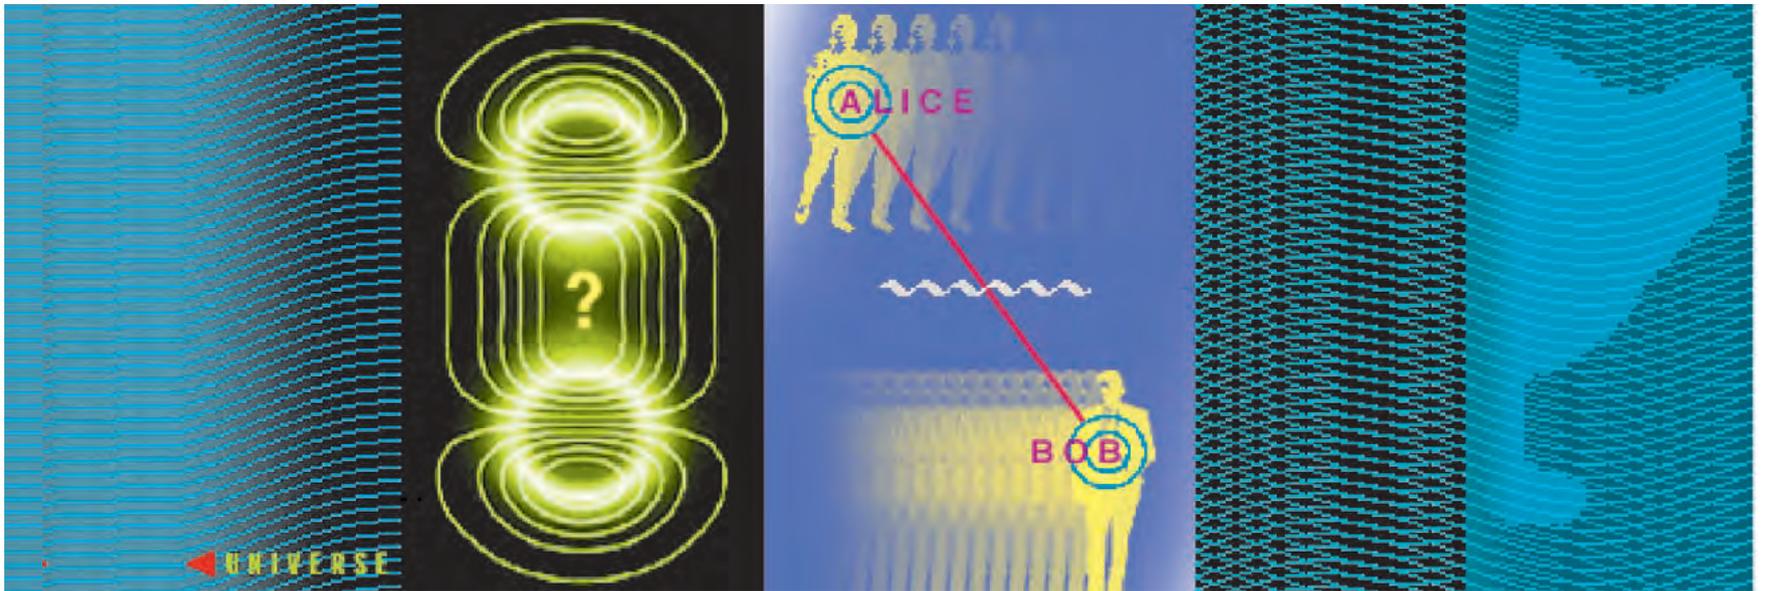
If only secrecy were so easy! To create a working quantum cryptography system, you first need to develop a tool to emit and selectively polarise individual photons. Photons – the smallest units of light – oscillate (or polarise) in various directions, but sending them through a polarising filter (think of Polaroid sunglasses) prevents certain orientations from getting through. So your system needs such a device at one end and at the other, a machine for measuring the polarisation of the incoming photons. That second job can be done with either another polarisation filter or, more commonly, by a birefringent crystal such as calcite, which sends photons into one of two paths depending on their orientation without absorbing them. The good news is that such instruments exist!

The big, it would seem insurmountable, problem is that in order to use quantum cryptography you'll need a direct connection between Bob and Alice, as using a router would have the same effect on the message as Eve attempting to read it. The current state of the art quantum cryptographic systems can send encrypted messages over a few kilometres.

## »» Oh. That seems a pretty huge limitation. Who'd want to use such a system?

Provided the problems with distance could be overcome, there are a number of ways a guaranteed secure line could be used. Banks and finance houses sending transactions across the globe, spies sending dispatches to their handlers, extremely rich philanderers planning illicit assignments, etc. If any

# WhatOnEarthQuantumComputing



readers discover a way of routing quantum communications, please let *LXF* know, in confidence, before you contact the patent office.

## » This is great! I was expecting numbers, equations and arcane formulae. Where's the mathematics?

Brace yourself. For the purposes of this limited exploration we'll keep the maths to a minimum because it gets very deep very quickly. We'll restrict ourselves to simple labelling, but if you want to know more, check out the resources below. Qubit labels look fairly similar to traditional binary: the quantum states attached to a single qubit are represented by  $|0\rangle$  and  $|1\rangle$  (qubit registers are shown in the same way e.g.  $|10101\rangle$  represents the number 21), and a superposition of these states is written thus:  $|0\rangle + |1\rangle$ . In terms of light, this equation tells us the light is polarised at 45 degrees, halfway between horizontal and vertical. To show light polarised at angles other than 45 degrees, we'd need to multiply each aspect by a 'weighting factor', and by varying this factor, an infinite number of superpositions can be represented. To complicate things further, if you're representing atomic spin instead of light polarisation, the possible angles are, of course, doubled as atoms are spinning in three dimensions.

## I want to know more!

### A little light reading

There just isn't enough space here to really get to grips with this subject, so try these resources. *The Fabric of Reality* by David Deutsch (Penguin Books, ISBN: 0140146903) takes you to the furthest reaches of scientific thinking, while *Minds, Machines, and the Multiverse: The Quest for the Quantum Computer* (Simon & Schuster Books, ISBN: 0684814811) by Julian Brown offers not just the theory behind quantum computing, but also the fascinating quest to make it a reality. *Schrodinger's Machines* by GJ Milburn (W H Freeman, ISBN: 0716731061) is also an excellent bedtime read.

This field moves pretty fast, so the web offers the

Right, that's all the time we have for numbers. Next question!

## » Forgive me if I sound a little impertinent, but what on "Earths" does all this have to do with either Linux or open source software?

While it's true that you're unlikely to see SuSE 15 for Quantum Computers in the near future, there are ways to simulate a quantum computer environment. And even though your simulations will never approach the power of a genuine quantum computer, it's still pretty cool.

You'd think that learning the methods for programming a quantum computer is pretty esoteric, and, you would be right. The always surprising Perl language has two modules available which simulate the behaviour of a quantum processes:

**Quantum::Superpositions** and

**Quantum::Entanglement**. Bearing in mind what we've already discussed, these modules do pretty much as expected.

**Quantum::Entanglement** was created by Alex Gough as an attempt to 'bring some of the functionality of the universe to Perl'. Variables begin in a superposition, but in the course of the programme, they collapse into a single value. When variables interact their fates become entwined,

best explanations of the current cutting edge. The first place to visit is [www.qubit.org](http://www.qubit.org) which has a selection of papers sorted by the knowledge required to read them. There's also a couple of good FAQs on that site. CPAN has all the details of the Perl modules mentioned above.

[www.dhushara.com/book/quantcos/aq/qcrypt.htm](http://www.dhushara.com/book/quantcos/aq/qcrypt.htm) has a good visual article which covers all the bases for quantum cryptography, and finally IBM's site about its own working quantum computer – [www.research.ibm.com/resources/news/20011219\\_quantum.shtml](http://www.research.ibm.com/resources/news/20011219_quantum.shtml) – demonstrates the current state of the art. It does look like wee you know.

meaning that observing one (and thus causing its collapse) will cause all linked variables to collapse at the moment of the observation. In order to simplify things, Gough says, the Entanglement module disregards 'eigen-functions, Hermitian numbers and other mathematical hurdles'.

The module adds an **Entanglement()** function which takes a list of amplitudes and returns a scalar in a superposition of values. You can create a superposition of 1 – 6 (to simulate the throwing of a die, for instance) with:

```
$die = entangle( 1=>1, 1=>2, 1=>3, 1=>4, 1=>5, 1=>6);
```

**\$die** will now, to all intents and purposes, contain each of those values until it is observed and collapses to a single state. To discover which way a variable collapses Gough says that as each value has "an associated probability value, all we need to do is build up a list of distinct outcomes, add up the amplitudes for each one, square the result, then use this to bias the value (or values) to which the variable collapses." So that's clear then.

The amplitudes in our example are all equal so:

```
print "You threw a $die.\n";
```

Will give you a one in six chance of rolling a number 1, 2, 3 etc. There are also options for complex numbers and variable amplitudes.

The **Quantum::Superpositions** module does a similar job but adds two new operators to Perl (any and all) which create a single scalar value (superposition) from a list of values (states). The any operator creates a 'disjunctive' superposition which can be in any one of its states at any given time, while all creates a 'conjunctive' superposition which is simultaneously in every one of its states.

## » Wait! What about Hilbert space, Fredkin gates, reversability, Grover's algorithm, Schrodinger's cat and Star Trek. Why haven't you told us about all that?

In fact I have mentioned all those things and somewhere in another universe, you have just finished reading about them. You enjoyed it too.

## Emulators



### Simon Goodwin rounds up of stray console emulators covers Playstation, Intellivision, Lynx and GameBoy Advance.



Your coverdisc holds a dozen freely-distributable Advance demos, including *FruitLand*, with source for the GCC cross-compiler, though its makefile presumes MSDOS paths, and Linux emulators for PlayStation, PC Engine, GameBoy Advance, Lynx, Chip8 and Intellivision.



Bally's Lynx game *Rampage* in a rotated, double-sized *Handy* emulator X window.

This is our last look at game console emulators in this series, taking in attempts to bring PlayStation, PC Engine and Intellivision software to Linux, plus emulators for Atari's Lynx and Nintendo's current GameBoy Advance. And if the Intellivision is not retro enough, try Chip8, an early cross-platform virtual machine which makes even that look high-resolution. Next month we return to home computers.

### Handy Atari

*Handy* is a free emulator for the Lynx colour hand-held computer, developed by former Amiga staff, funded by Epyx and eventually marketed by Atari. Good games, stereo sound and a colour LCD were Lynx strengths, but meagre marketing and battery life limited sales, though they're no obstacle to emulation.

*Handy* comes with HTML documentation and automatically unpacks 'zip' and 'gzip' ROM files. It uses *SDL*, the *Simple Direct media Layer*, for hardware abstraction. It supports sound, scalable 16-bit colour graphics, joysticks, screenshots and game saves.

Your X display must be in a 16-bit colour mode. *SDL 1.2.2* or later is recommended, and the *libz* compression library which is standard on most Linux distros. The emulator comes precompiled for recent Linux systems – it called for *libstdc++-libc6.2-2.so.3*, which is not part of the Debian 2.2 distro, but a link to *libc6.1-1-2-1.9.0.so* got *Handy 0.82R1* running.

Graphics are fast, accurate and system-friendly, though sound was rather jerky with the screen scaled up to the maximum X window size, using 16 pixels for each original. Sound and vision were smooth and larger, though still only filling about a fifth of my monitor, in *SDL's -fullscreen* mode.

Besides game cart images you need a copy of the small Lynx ROM bootstrap, which is not bundled for copyright reasons. Many games expect four joystick or pad buttons in addition to directional control, though you can use **Z**, **X**, **1**, **2** instead, with cursor arrows on a keyboard, plus **F10** to reset and **Enter** to pause. The Lynx facility to rotate the screen left or right, vital for games like *Klax* and *Gauntlet*, is triggered with **F11** and **F12**, while the **Home** key resets the normal view. Initial **-scale**

parameters from 1 to 4 adjust the emulated display window size. **F3** and **F4** save and restore game state, **F9** saves up to ten snapshots of the screen, and **Esc** quits the emulator.

### Intelligent TV?

Mattel's Intellivision console provided stiff competition for Atari's VCS two decades ago. The CP1610 central chip is a true 16-bit processor though its CISC instructions are assembled from ten bit sections. Sound comes from a three channel GI beeper later popularised in Spectrum 128 and Amstrad CPC home computers.

There are two Intellivision emulators for Linux. *Bliss* is best for the casual x86 gamer and former user, while *jzIntv* gives greater insight into the depths of Mattel's machine and is more portable.

You need ROMs from the original system, in addition to game cartridge ROMs: the small graphics *grom.bin* and the kernel *exec.bin*, which loads at \$1000, with games at \$5000. Some games also require *ecs.rom*, and Intellivoice speech synthesis demands *ivoice.rom*.

### Bliss

*Bliss* comes in versions for Java and X, with a perfunctory Readme file. The X version, built on *SDL 1.2.3* and *Zlib 1.1.3*, is sluggish so the Java one will try the patience of all but the catatonic or bleeding edge Virtual Machinists. At first *BlissX* sniffed at my Debian 2.2 system because it required



Intellivision *Dungeons and Dragons* zoomed up on X by *Bliss*.

*libSDL\_image-1.2.so.0*, absent even after building *SDL 1.2.3* from source. *libSDL\_image* is an uncommon requirement, packaged apart from the main *SDL* binary and source collections; it's not in the stable Debian distro, but will be in 3.0. *Bliss* key bindings are fixed, with digits for the calculator grid of the original Mattel keypad, **Del** for clear and **Z**, **Ctrl** and **X** for the three action buttons. **Enter** and the arrow keys complete the control set.

*jzIntv*, Joe Zbiciak's *Intellivision Emulator* – better documented than *Bliss*, and less fussy about *SDL*, though linking may require

makefile tweaks to suit your library locations. It's released under GPL and also suits MacOS, so it's not tied to Intel architecture.

The documentation includes a well-written CP-1600 programming guide, aimed at people who know nothing more than a little BASIC. The overview of the source and utilities is particularly helpful, though rather old. You get utilities to convert ROM images between 10 and 16-bit format and a CP-1600 code disassembler. The ESD-based sound is CPU-intensive but accurate, mixing 16-bit data at sample rates from 4 to 48KHz.

## GameBoy Advance

**N**intendo's GameBoy Advance is a modern hand-held console running 240 by 160 pixel 2D colour graphics, 384K of RAM in three sections, and games on ROM cartridges, typically a few megabytes each. The British-designed ARM7 RISC processor is clocked at 16MHz, fast enough to offer interesting gameplay without outflanking interpreting emulators on a modern desktop Linux system.

The Advance offers four sampled sound DMA channels as well as GameBoy beeps. The display supports sprites with 16 or 256 colours, hardware scrolling, tile flipping and rotation, and a high colour display, souped up from the Super Nintendo console.

An eight-bit Z80 coprocessor lets it run older GameBoy software. My *LXF21* column listed ten emulators for such games and demos, so GBA emulation would be the long way round, but the new handheld has been adopted fast by emulator writers as well as gamers; there are already three competent Advance emulators for Linux.

It's feasible to develop Advance software on your own, making this a good route into the games industry. High ROM prices and low margins leave a niche for freelance programmers familiar with classic games and *GCC*. Demo sources, tools and examples abound, though you'll need backing from a Nintendo licensee to get your game into the shops.

## Boycott Advance

*Boycott Advance* is beerware from Niels Wagenaar of The Netherlands, who'd appreciate a postcard if you can't send beer. The release version is not crippled, so registration just gets you (and Neils) a warm feeling, access to pre-releases and your name in the credits.

I tested version 0.22 of the *SDL* remix. Base requirements are a 500MHz x86 with 64MB RAM and 16-bit colour. It managed about 45 frames per second on the demo game, a neat revamp of *Pong*, in the default tiny X window. This rate falls to 15 to 20 in the top scale setting, but remains playable with a frame skip of 3.

Emulation of the 32-bit ARM processor, rather than the sound or graphics, makes Game Boy Advance emulators greedy for CPU time, but the X overhead is substantial; after selecting **-fullscreen**, *Boycott* delivered smooth graphics at 70 Hertz, outrunning the nominal 60 Hertz LCD, on a custom console the same size as the **-scale 4** window. The switches must follow the command and ROM name, peculiarly.

Four adjacent keys, **Z**, **X**, **C** and **V** mimic left, A, B and right buttons respectively, with **Enter** to start, **Tab** for select, **F10** to reset and **Esc** to quit. Arrows, joy pad or joystick indicate directions, while **F9** grabs a screenshot, **F1** pauses and **F2** continues. **F3** and **F4** control the frame rate display in the

window border, and **F5** and **F6** trim the update speed, though they respond so fast *Boycott* often skips from one in eight down to 1:1 or back again each time you press the key, requiring several tries to chance upon any intermediate setting.

## VGBA

*VGBA* is the latest from Marat Fayzulin, author of fine MSX, NES and earlier GameBoy emulators. I tested the pre-compiled version 1.2 which unpacked and ran readily on my Debian 2.2 system.

*VGBA* comes with HTML documentation cloned from Marat's other emulator docs. The recommended minimum host is a 400 MHz PII with fast graphics. **-sync** can limit fast machines to accurate speed by tying the emulator to your screen refresh rate.

You must run your X display in a 16-bit mode as the console uses 32K colours. Version 1.2 is a major update with much better sound and interrupt handling than previous *VGBAs*. Prototypes ran on MSDOS, but lacked many features. The Win32 version is commercial so *VGBA* is not open source; ports are available for several Unix architectures, but the code relies on a little-endian host CPU, ruling out versions for SPARC, 68K or PPC. The sound driver uses Posix threads, standard on Linux.

*VGBA* runs most commercial GameBoy Advance ROMs, but mainly addresses game developers and demo programmers who can't afford Nintendo's devkit but want to impress potential publishers. To avoid copyright problems the Nintendo BIOS is emulated by trapping SWI instructions. This is less hardware-compatible but faster than running the original BIOS, emulating right down to the silicon. It's only likely to fail on demos that break Nintendo development guidelines by poking the hardware directly.

The main control keys are **Space**, left **Ctrl**, **Tab**, and **Enter**, but sundry letters are polled simultaneously so it's easy to find an ergonomic layout. Function keys save and load state in '.STA' files, toggle autofire on the four main buttons, reset with **F11** and quit with **Esc** or **F12**. You must quit to load a different ROM.



**Blues Brothers** for the GameBoy Advance emulated courtesy of Titus and *VGBA*.



GameBoy example *FruitLand*, shakily rendered by *Boycott Advance*.



# Emulators

## Chip8/vision8

### Pioneer

*Vision8* emulates a pioneering cross-platform virtual machine. *Chip8* was introduced in 1975, on RCA CDP1802-based computers, like the RCA's **TELEMAC 1800** home console, bundled with games on cassette tape, and the **COSMAC VIP**, advertised to UK hobbyists as a home computer into the 80s.

*Chip8* games were played with a hex keypad or joystick, using chunky mono graphics reminiscent of early TV games. Besides various puzzle games, *Pong* and *Brix* revisit *TeleTennis* and *Breakout*, each in under 300 bytes, with on-screen scoring. *UFO* is slimline *Space Invaders* with a base.

*Vision8* runs on Unix systems with X in 8, 16 or 32-bit per pixel modes, MSDOS with CGA, 8 bit MSX, ColecoVison and Adam micros, and 32-bit AmigaOS 3 systems. Its portability is exemplified by recent ports to GameBoy Advance and HP48 calculators! I renamed the variable **sync** in *X.c* to get *Vision8* to compile with *GCC*. Renaming it **synch** in three places resolved a clash with the GNU function that ensures disk contents are flushed to drive hardware.

*Vision8* is small and elegant, though the appeal is minimalist. With its handful of opcodes and tiny programs, the gameplay to program size ratio beats many modern CD and DVD releases.

◀ **F1** calls up the debugger and numeric keys toggle sprites, 4 background graphic layers, and six sound channels. Audio volume is also key-controllable, though there's no GUI or menu system.

The nicest feature is the way you can stretch the screen to almost any size by dragging the corner of the window as the emulator runs – after a moment the display will resize to fill the new window, smoothly scaled to fit, with no obvious stepping of pixels. It even allows views to be scaled down from the normal GBA resolution, but only steroidal WAP telephonists are likely to benefit from that. Substantially stretched windows grew slow on my K6/500 unless I used **-uperiod** to skip screen updates.

*VisualBoyAdvance* also runs programs for earlier GameBoys, like the real thing unlike other Advance emulators. It uses SDL on Linux and Windows, requiring a 500 MHz PIII for full speed. My AMD K6 delivered playable results, but sound was uneven and it struggled to reach 90% performance on Advance emulation. Plain GameBoy and GBC support is far less demanding.

## VisualBoyAdvance

I tested version 0.8, supplied as an x86 binary with a small, well-commented configuration file, and brief readme text. Numeric pad keys are said to indicate directions but I found that cursor arrows worked, in practice.

Keys are definable by changing the SDL keycodes in the '.cfg' file. This can also select digital joysticks or pads, one to four times zoom or full screen display, half a dozen video filters, and frame skipping in Advance and standard GameBoy mode, from 0 to 5



*Lego Racers 2* on *VisualBoyAdvance*.

frames. MMX instructions and a real GameBoy BIOS may be switched in, but are not required. Border and status displays are also configurable, along with the paths used for ROM loads, screen and state saves, and storage of 'battery-backed' data.

Function keys 1 to 10 save game states, with shift to reload, so you can save and restore ten positions without an external script to shuffle multiple saves. **F12** grabs a screen. **Alt-1** to **4** toggles auto-repeat on the A, B, left and right buttons, mapped to **Z**, **X**, **A** and **Z** (again?!). – at least according to the README, the fourth one should say **S**.

**Enter** mimics Start and **Backspace** is Select. **Ctrl-R** resets and **Ctrl-P** pauses the emulator. **Esc** quits and **F11** enters the debugger, undocumented but with adequate online help. This offers ARM and Thumb code disassembly, breakpoints on access or write, single-stepping and facilities to view and edit memory, ARM and I/O registers.

## Comparisons

*Boycott Advance* runs some ROMs that *VGBA* rejects or renders blankly, like the mode 7 Doom scrolling demo, but *VGBA*'s **-nocrc** switch reprieves some with incorrect ROM checksums – another area where *VGBA* sets out to be strict. *Boycott*'s **-romwrite** switch gets some demos running that would fail on a real Advance. *Boycott* suffers from graphics glitches, especially on text, and *VisualBoy Advance* could do with better sound and documentation, but all these emulators work well on reasonably modern PC hardware.

## PC Engine emulators

### Long-running 16-bit consoles

PC Engine 16-bit game consoles, also branded TurboGraphX, were made in Japan by NEC until the early nineties, strongly supported by Hudson, now Nintendo partners. Mostly rôle-playing and shooting games were produced for it, from 1985-99. *VPCE* and *Hugo* are the two PC Engine emulators for Linux. Both work well, after tweaks, but *VPCE* gives most insight into the Engine.

*Hugo* uses *Allegro* for platform-independent sound and graphics. The docs are aimed at Windows users, but have a brief Linux section. A bundled *Pong* game helps you prove that it

works, but installation is tricky even though the x86 code comes ready-built. The archive has incorrect permissions so the executable needs 'chmod' to bring it to life. You also need to copy a data file: `mv hu-go!.dat /etc/hugo.dat`

*VPCE*, the *Virtual PC Engine*, comes pre-compiled as Intel code for *libc5* and *glibc2*. Each archive contains a small readme file and executables that run in a X window or full screen via *SVGAlib*.

To run the *glibc2* version on my Debian `ln -s /usr/lib/termcap.so /usr/lib/termcap.so.2`

That made a link from the installed

terminal capabilities to the specific version requested by *VPCE*. Other versions bombed with a Segfault. Games are not bundled but *VPCE* unzips ROMs automatically, & runs *Hugo's Pong* demo.

The normal screen mode uses 256 by 240 pixels at 60 Hertz, but **-mag n** zooms the screen **n** times in X and Y and **-frameskip** skips frames for slow hosts. *VPCE* defaults to 8-bit colour but shifts to 16-bit if that's your XFree86 setting; otherwise **-truecolor** sets other depths.

Sound uses '/dev/audio' for 8-bit µlaw, unless you specify **-dsp**. As with *Hugo*, I found it rather grainy. Move around with

cursors, a joypad or joystick, via '/dev/js0': **C**, **X**, **Return** and **Space** keys emulate console buttons, and **Alt-F3** saves a TGA screenshot. Press **Esc** to quit, twice if in **-debug** mode.

The **-debug** command option invokes a fine debugger. A VRAM browser identifies background, map, palette and sprite graphics. The monitor terminal can disassemble code for the 68516 processor, as used in Apple GS and SNES systems. Digit keys **1** to **4** generate interrupts. You can step, skip, insert **NOPS**, **RETs** and **RETIs**, review memory or enter 'R' to run.

## PlayStation

Sony's PlayStation (PSX) game system needs little introduction. It's a CD-based console with a 33MHz MIPS R3000 processor, SPU audio engine, and capable though obsolescent 3D graphics. It has sold tens of millions and still thrives, eight years on.

There are six PSX emulators for Linux; most share a plugin system, so the emulator contributes the GUI, CPU and memory emulation core, and you mix and match add-on code for sound, graphics, CD, pad and network support to suit your host. Emulators and plugins are typically available for Windows and x86 Linux systems, so the name alone does not guarantee a GNU-friendly binary. All the emulators rely on a copy of Sony's ROM BIOS from a real PSX; the US version *scph1001.bin* gives best compatibility.

DIY PlayStation software development is now feasible using open source software plus an *Action Replay* or similar cheat cartridge, patched BIOS, and an ISA card with a bi-directional parallel port. Our links point to penguin-friendly PSX development sites; emulators are a big part of this scene.

### PSX Emulators

*PCSX* development is Windows-led, but the Linux version is at 1.2 and capable, if a little flaky. Console messages warn of GUI bugs while setting up and it rejected some the plugins I tried.

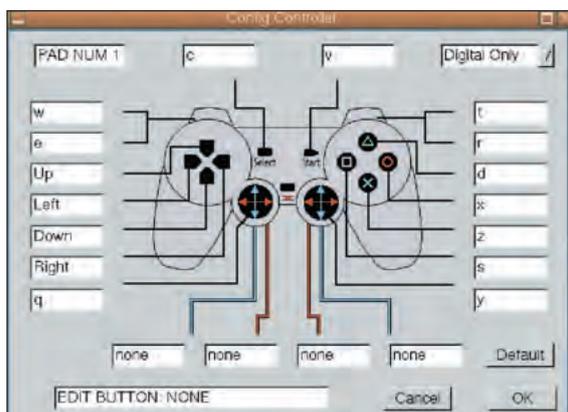
*ePSXe* also runs on Windows, but the Linux version is slick, solid and feature-rich. Sadly the binary is Intel-specific, whereas *FPSE* runs on Windows and PowerPC AmigaOS, so it is not hard-wired for little-endian systems, though BeOS and Linux versions have vanished from the home page. *GPSE* is a fork of *FPSE* 0.08 made under GPL, with the agreement of *FPSE*'s main author.

*SOPE* development stalled late in 2000. Version 0.04 was released under GPL and has dynamic compilers for Alpha and x86 and interpreters for other systems. You may encounter two other contenders, though work on them seems to have stopped. *PEX* was developed in 2000 with Slackware, but the site was down each time I tried – the emulator seems to have vanished.

### PSX tests

My tests used Pete's *SoftX driver1.52* and the *Peops* audio driver 1.1 for OSS. Audio has a configuration GUI, but graphics options are set in a '.cfg' file.

I found cover-mount game demos a cheap and effective test



This *ePSEe* dialogue maps PlayStation buttons to Linux keys.

of compatibility. *ePSXe* was much the most capable running games like *Driver*, *Gran Turismo 2*, *TOCA*, *Tomb Raider TLR*, *International Track and Field* and *FIFA 2000* playably, though at frame rates between 10 and 20 rather than the 50 or 60 they were made for. 3D hardware and a modern CPU will close this gap – if you already play 21<sup>st</sup> century action games, your system should cope fine. The only unplayable game was *Metal Gear Solid*, which flickered to the point of uselessness on my setup.

Streaming CD audio sounded choppy, but in-game effects sounded fine. The graphics plugin saves screens in bloated BMP format, at 1.5MB a time, whereas *XV* packs the same images well within 100KB as PNGs.

### PS2 prognosis

Sony have since moved on to the 128-bit PlayStation2 (PS2). You're not going to emulate the PS2 on any desktop computer for a few years yet, but *Pcsx2* might give useful insights into the architecture. A more practical option is Linux for the PS2. [LXF](#)

**Konami International Track and Field literally runs on ePSXe.**



### Links

Advance demos: [www.gbadev.org](http://www.gbadev.org)  
 BoyCott Advance:  
<http://boycottadvance.emuunlim.com/>  
 Emu series online:  
[www.simon.mooli.org.uk/LXF](http://www.simon.mooli.org.uk/LXF)  
*EPSXE* home: [www.epsxe.com](http://www.epsxe.com)  
*FPSE* home: <http://fpse.emuunlim.com/>  
*GPSE* CVS:  
<http://savannah.gnu.org/projects/gpse/>  
 Hugo home: *FPSE* home:  
<http://emuunlim.com/hugo/>  
 Intellivision extras: [www.intellivisionlives.com](http://www.intellivisionlives.com)

*PCSX* home: [www.pcsx.net](http://www.pcsx.net)  
*PSX* emu UK:  
<http://linux.psxfanatics.com/>  
*PSX* Tux devkit: [www.psxdev.de](http://www.psxdev.de)  
 Simple Direct Media Layer: [www.libsdl.org](http://www.libsdl.org)  
*SOPE* home: <http://sope.sourceforge.net/>  
*VGBA* home: [www.komkon.org/fms/VGBA](http://www.komkon.org/fms/VGBA)  
*Vision8* home:  
[www.komkon.org/~dekoegel/vision8.html](http://www.komkon.org/~dekoegel/vision8.html)  
 VisualBoy Advance: <http://vboy.emuhq.com/>  
*VPCE* home:  
[www.geocities.com/SiliconValley/Way/3340](http://www.geocities.com/SiliconValley/Way/3340)

# 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 website 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...



### Intel C++

Speed up your C, C++ and Fortran code with Intel's 32 and 64-bit compilers – find them on your coverdiscs **p68**

### Java

Textual search and replace operations using regular expressions and patterns for our mini-project **p78**

### Gnumeric

Our practical tutorial takes a look at using GNOME's spreadsheet to get some real work done with your Linux office suite **p70**

### Kylix

We apply encapsulation, inheritance and polymorphism to our code as we introduce OO principles with Object Pascal **p80**

### Perl

Using the popular Tk toolkit we help you to get graphical and add some widgets to your Perl programs **p74**

### PHP

Our new series introduces the popular and effective web scripting language – starting you on dynamic web content **p86**

## TIP OF THE MONTH!

Linux doesn't impose the stupid restriction that Windows does that a file's type must be signified by a three letter filename suffix. Chalk one up for flexibility, but how is one supposed to easily tell a file's type? Simple: the *file* command.

The *file* command, present on all Linux distros, determines a file's type by examining its content. First it checks whether the file in question is a special file, such as a device or a pipe. If not, it then uses a database of rules or so-called magic numbers – usually stored in the file

`/usr/share/misc/magic` – to determine it's type. If no rule matches, it checks to see if the file matches any of the popular encodings used for text files.

Using *file* is easy. To query the file 'somefile' in the current directory:

```
file somefile
```

If 'somefile' is compressed with *gzip*, then *file* will return telling you so. To examine the content of a *gzip'd* file use the `-z` switch:

```
file -z somefile.gz
```

Alternatively, you may pipe the file to be examined as follows:

```
zcat somefile.gz | file -
```

This is particularly useful for *bzip'd* files, because *file* cannot decompress them itself:

```
bzcat anotherfile.bz2 | file -
```

The *file* command has rules to determine filesystem types, not just data files. This can be used on files containing filesystem images – or actually on devices, if you use the `-s` switch. *E.g.* to determine the filesystem stored on `/dev/hda3`

```
file -sL /dev/hda3
```

The `-L` switch here tells *file* to follow any symbolic links – handy if you are using devfs (and `/dev/hda3` is really a symbolic link).

## 32/64-BIT COMPILERS

# Using Intel's C and C++ Compiler

**PART 1** Intel's Getting Started Guide will help you to get the most from the compilers on our coverdiscs.



This tutorial explains how to install the Intel C++ compiler for Linux, build "Hello World" for IA-32 and Intel Itanium™ processor based systems and how to get started optimising your applications with the Intel compilers. The Intel C++ Compiler 6.0 for Linux consists of the following:

- Intel C++ Compiler for IA-32 based applications: icc
- Intel C++ Compiler for Itanium processor-based apps: ecc
- Linux Application Debugger 6.0 for IA-32 based apps: ldb
- Linux Application Debugger 6.0 for Itanium based apps: ldb
- Intel Itanium Assembler 6.0 for Itanium based applications: ias
- Product documentation

## Installation Notes

The Intel compilers use the GlobeTrotter *FLEXIm* electronic licensing technology. A valid licence is needed to use these compilers. To install the compilers:

- 1 Insert and mount the *Linux Format*

coverdisc. Open a shell, *cd* to a temporary directory, and copy the compiler package `l_cc_p_6.0.139.tar` there from Magazine/Intel.

- 2 Unpack the compiler package with the command

```
tar xvf l_cc_p_6.0.139.tar
```

- 3 Become the root user, and execute the install scrip with

```
./install
```

It's possible to install without root access by unpacking the RPM files with *rpm2cpio* and editing the compiler environment and configuration files, but this installation method is not supported.

- 4 Select the compiler or debugger to install. The default RPM options **-U --replacefiles** are recommended to update existing files. The recommended installation directory is `/opt/intel`.

- 5 After installation, press **x** to exit.

- 6 Execute the appropriate script to setup the compiler environment variables (assuming you installed in `/opt/intel`). For the IA-32 compiler enter:

```
source /opt/intel/compiler60/ia32/bin/iccvars.sh
```

For the Itanium compiler enter:

```
source /opt/intel/compiler60/ia64/bin/eccvars.sh
```

(For those using *csh*, the appropriate scripts have the file suffix `.csh`.)

- 7 Install the *FLEXIm* license. You need to register the compiler online at

<http://www.intel.com/software/products/distributors/linuxenterprise.htm>

- 8 Run the compiler. The compilers use environment variables for configuration, and the install script creates script files to correctly set the compiler environment. The names and locations are

```
/opt/intel/compiler60/ia32/bin/iccvars.sh
```

for IA-32 applications, and

```
/opt/intel/compiler60/ia64/bin/eccvars.sh
```

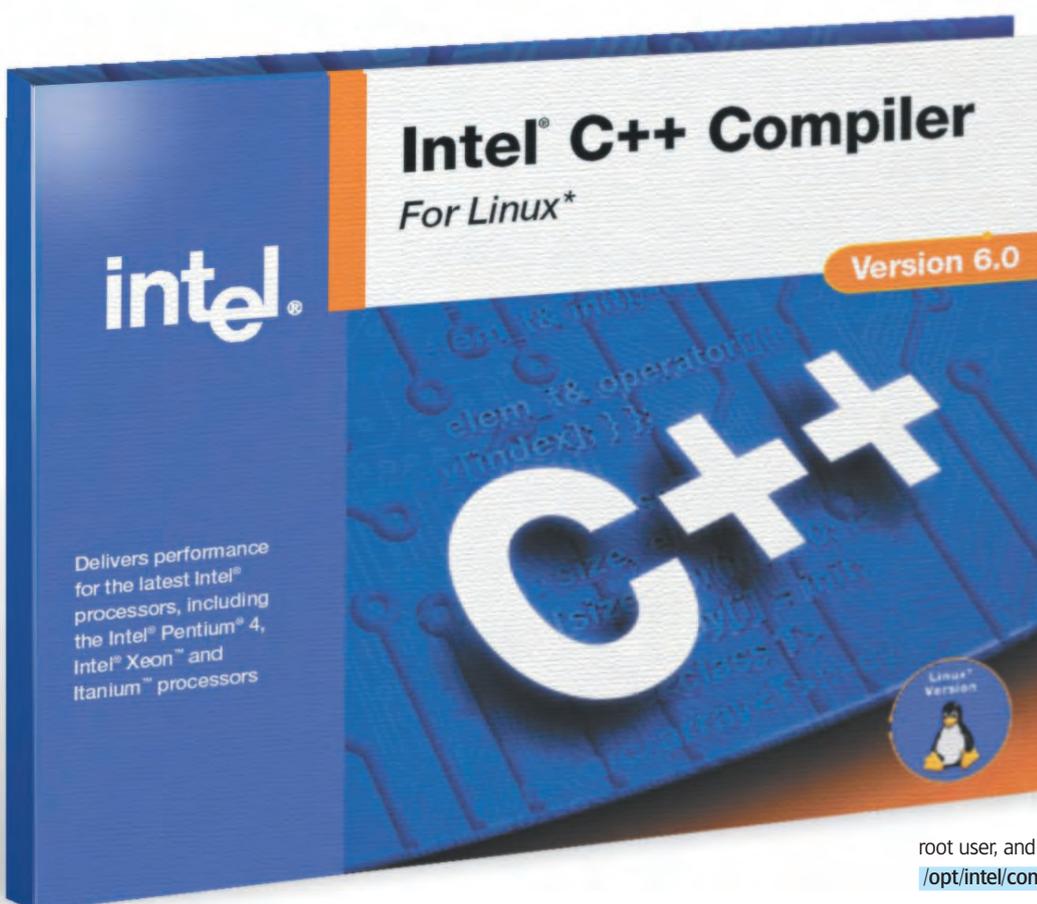
for Itanium processor-based applications.

The installation creates configuration files, `/opt/intel/compiler60/ia32/bin/icc.cfg` and `/opt/intel/compiler60/ia64/bin/ecc.cfg`. These contain common settings and can be used to add additional default options. Note, if you install a compiler update package, you need to rename the config files to avoid being overwritten.

**To uninstall the Compiler:** Become the

root user, and run the uninstall script:

```
/opt/intel/compiler60/ia32/bin/uninstall
```



## System Requirements

### IA-32 Processor System Requirements

- |   |   |
|---|---|
| <ul style="list-style-type: none"> <li>■ A computer based on a Pentium processor or subsequent IA-32 based processor (a Pentium 4 processor is recommended)</li> <li>■ 128 MB of RAM (256 MB recommended)</li> <li>■ 100 MB of disk space</li> <li>■ A supported Linux distribution (see <i>Supported Linux Distributions box</i>)</li> </ul> | <p><i>Linux Distributions box</i>)</p> <p><b>Itanium Processor System Requirements:</b></p> <ul style="list-style-type: none"> <li>■ A computer with an Itanium processor</li> <li>■ 256 MB of RAM</li> <li>■ 100 MB of disk space</li> <li>■ A supported Linux distribution (see <i>Supported Linux Distributions box</i>).</li> </ul> |
|---|---|

## Supported Linux Distributions

This version of the Intel compilers for Linux supports Linux distributions with the following combinations of glibc, the C language library, and the Linux kernel. Supported IA-32 based systems:

- Distribution with 2.2.2 glibc & 2.4 kernel; or
- Distribution with 2.2.4 glibc & 2.4 kernel.

Supported Itanium processor based systems:

- Distribution with 2.2.3 glibc & 2.4 kernel; or
- Distribution with 2.2.4 glibc & 2.4 kernel.

on IA-32 based systems, and

`/opt/intel/compiler60/ia64/bin/uninstall`

on Itanium-based systems.

## Building “Hello World”

Building the classic “Hello World” program is described. After successful installation with a valid *FLEXIm* license, the configuration scripts, *iccvars.sh* or *eccvars.sh*, must be sourced as described above. Then, execute these steps:

**1** Create a simple “Hello World” C++ program in a text editor and call it ‘hello.C’:

```
#include <iostream>
int main()
{ std::cout << "Hello World !" << std::endl; return 0; }
```

**2** Set up the Intel C++ Compiler environment variables as discussed above: *iccvars.sh* for IA-32 applications or *eccvars.sh* for Itanium-based applications.

**3** Compile hello.C. For IA-32 applications, enter:

```
icc hello.C -o hello
```

For IA-64 applications, enter:

```
ecc hello.C -o hello
```

**4** Run the executable with

```
./hello
```

It should display “Hello World!”.

## Compiler Optimisations

The Intel C++ Compiler enables programmers to take full advantage of the advanced performance enhancement features of Intel’s latest IA-32 and Itanium processors and includes advanced optimisations. These optimisations are intended for use in product-release builds of applications, not necessarily for earlier phases of application development cycles. In general, increasing the degree of optimisation done by the compiler leads to an increase in compile-time and reduced debugging capability. This section describes an optimisation methodology with the Intel C++ Compiler.

During the application development, the **-g -O0** switches are recommended to allow fast compile times and full debugging with no optimisation. To start to optimize, the default optimisation **-O2** is recommended. The **-O3** enables advanced optimisations. Interprocedural optimisation allows the compiler to optimise across different compilation units and can have large performance improvements. Profile guided optimisation uses information from running an instrumented executable that allows the compiler to rebuild the application knowing where the majority of the computations are. Of course, not all optimisations are beneficial for all applications.

For additional details on optimising, the paper, *Optimizing Applications with the Intel C++ and Fortran Compilers*, is available at [www.intel.com/software/products/compilers/c60/linux](http://www.intel.com/software/products/compilers/c60/linux). For complete information on the individual optimizations, please refer to the installed Intel C++ Compiler User’s Guide at `/opt/linux/compiler60/docs/c_ug_Inx.pdf`.

Remember to always measure the performance of your application after each optimization added to verify the benefits.

## Additional Information

Your feedback is very important to Intel. To receive technical support for the tools provided in this product and technical information, including FAQs and compiler updates, you need to be registered for an Intel Premier Support account at <https://premier.intel.com/>. Note that, if you already have access to Intel Premier Support from another Intel product, then you do not need to re-register. Compiler support information, including top technical issues and known issues, is available at <http://support.intel.com/support/performance/c>. The product release notes contain additional information and can be found at `/opt/intel/compiler60/docs/C++ReleaseNotes.htm`. The Compiler User’s Guide is at `/opt/intel/compiler60/docs/c_ug_Inx.pdf`. Have fun compiling! [LXF](#)

## The benefit of Intel’s compilers

### GCC’s rival puts forward its case

Special optimisations for P IV, Celeron, Xeon and Itanium processors make it easy to get outstanding performance for your applications.

The Intel C++ compiler supports the new 3.0 ABI standard – but also works hard to ensure code integrity with files compiled on GCC.

As well as profile-guided optimisation and interprocedural optimisation (IPO), the compilers feature

data-prefetching and floating point optimisations for significant performance improvements in computer intensive applications.

IPO can dramatically improve performance in programs that contain many small or medium sized functions that are frequently used, especially for programs that contain calls within loops.

Also featured is support for the P IV’s hyperthreading.

## Copyright and Legal Information

Intel, Itanium, Pentium, and VTune are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries

Other names and brands may be claimed as the property of others

Copyright © 2001, 2002 Intel Corporation. All Rights Reserved.



## USING LINUX SPREADSHEETS

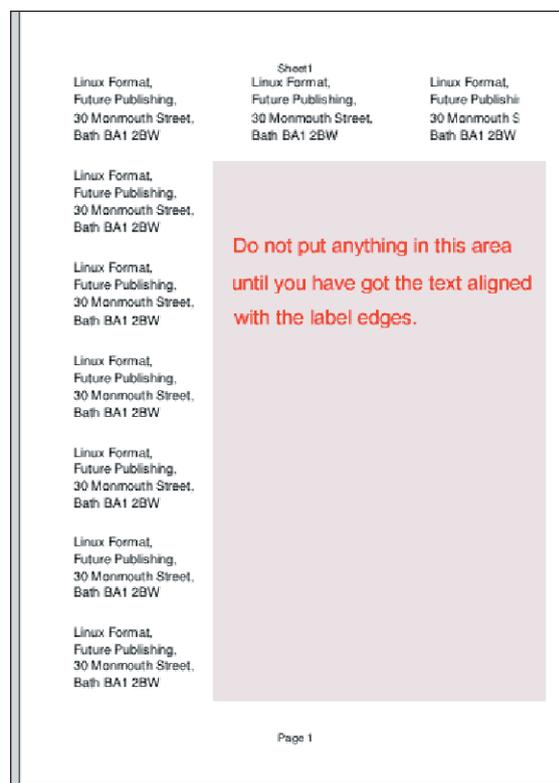
# Getting organised with Gnumeric

Neil Lucock gets into Gnumeric and shows how you can keep track of things with spread sheets.

So, what can you do with this spreadsheet thing that you found hiding in the menu under "Office"? There are two provided in the distribution I use, *Gnumeric* and *KSpread*. Once you can understand what you can do with one, using the other should not be a problem. I intend to use *Gnumeric* to make something useful, but the procedures will apply to any spreadsheet program you need to use. I'll assume that you can get the thing to run, know how to open and save files and have a basic familiarity with some spreadsheet functions. This article was written using *Gnumeric 1.0.4* on Mandrake 8.2 in KDE. Older versions may not have all the functions.

Some people just use a spreadsheet as a layout tool. If you have sheets of sticky labels, you can use *Gnumeric* to print out the address on each. You can type your address into the cells, copy and paste it across the width of a page and down one edge, then do a test print on plain paper. There's no point in wasting ink, print the full sheet when you have got it right. See where the print comes out when you compare it to your label edges, adjust it by moving the gridlines (the edges of the cells and columns) so that it will align correctly when you print on the sheet of blank labels (this might take a few goes to get it right). Once it is aligned, copy and paste the address into the rest of the sheet. Once you have made one that works, save it to use as a template. After that changing the address is just a copy and paste job. Ensure the font characteristics remain the same as the original, though. The "print preview" is a useful tool to get it nearly right.

Right: Using a spreadsheet to print labels. Perhaps not the best use of *Gnumeric's* abilities.



- Applix (\*.as)
- Data Interchange Format (\*.dif)
- EXPERIMENTAL SAX based Gnumeric (\*.gnumeric)
- GNU Oleo (\*.oleo)
- HTML file made by gnumeric
- Linear and integer program (\*.mps) file format
- Lotus 123 (\*.wk1)
- MS Excel (tm) (\*.xls)
- MultiPlan (SYLK) import
- Plan Perfect Format (PLN) import
- SC/xspread
- Text File import (customizable)
- Xbase (\*.dbf) file format

**Gnumeric's import options.**

If all you ever do with a spreadsheet is use it as a simple layout tool, you are missing out on a lot of useful functions. *Gnumeric* competes with Microsoft's *Excel*, which, like it or not, is the standard. If you use *Excel*, most of *Gnumeric's* icons and functions will be familiar to you. Businesses want to do the same things, no matter whose software they use. No one is going to change to Open Source software if the benefits are instantly wiped out by a loss of productivity as staff learn unfamiliar icons and commands.

*Gnumeric's* import filters can read a wide variety of file types, including *Excel*, Lotus and Applix formats. The original idea for this spreadsheet came from an *Excel* spreadsheet used by Southern Graphics. It was emailed to me and *Gnumeric* opened it with no problems. However, it was comparatively simple. If it includes drawings done with the *MS Office* drawing applet, it might not import correctly.

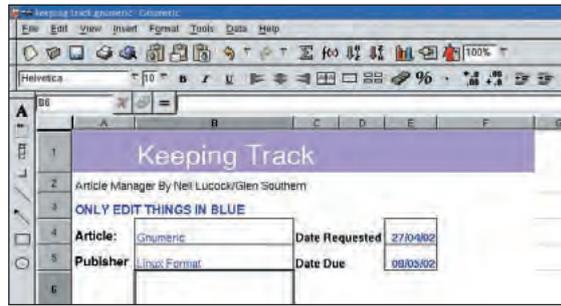
**Creation**

The first step in creating any spreadsheet is to sit down with a coffee and decide what you want it to do. What I want is something to keep track of how my *Linux Format* articles are doing. The shoebox with loads of scraps of paper is getting full, so I'll need something to show exactly where I am and where I should be as the publisher's deadline approaches.

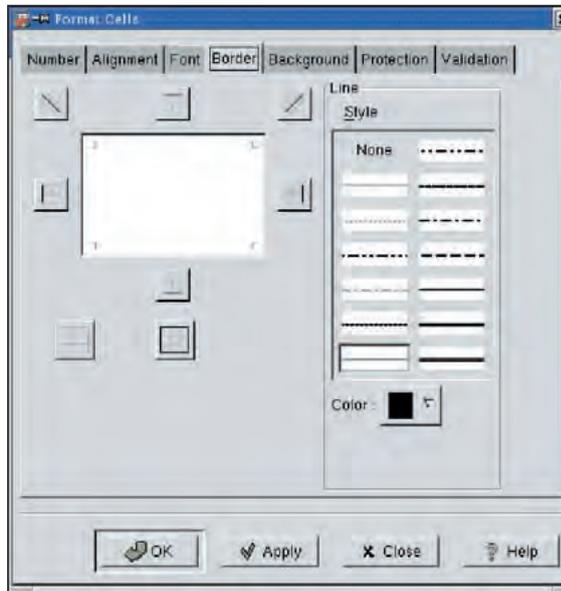
I've started with a new spreadsheet and put in the title and a reminder to only change values in cells with blue text. I've put in labels for the article name, the publisher and the dates articles were commissioned and when they are due. I've got the cell alongside "Date Requested" selected.

If you choose Format>cells> and click the "border" tab, you can put a border around the cell to show that something should go there. Clicking the "font" tab allows us to set the default colour for the selected cell. I've made it blue. When I type something in the cell and press "enter", the contents will become blue. With the cell still highlighted, also click the "Number" tab and make the cell type "date" and select the format you prefer. The cells for Article and Publisher are formatted as "Text" and made blue with borders. The Validation tab can check the type of data typed in and display an error message if you mistakenly type in a name instead of a date, for example.

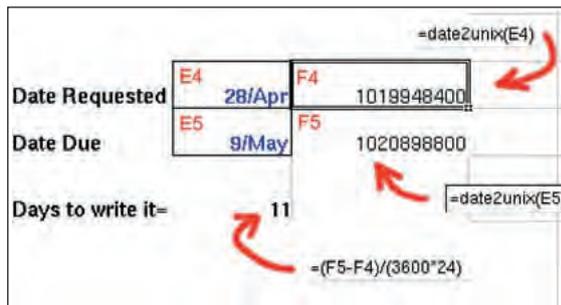
I can put the current date in to a cell by using **Ctrl-;** (Control and semi-colon) or use "Insert>Special>Current date" from the menu. The email or letter from the editor will say when the



The new sheet takes shape with a reminder to only edit blue things.



The Format>cells tabbed dialogue for Borders.



The DATE2UNIX formula.

Section	Details	Words	Total
Introduction		115	
Subject 1			
Subject 2			
Subject 3			
Subject 4			
Subject 5			
Subject 6			
Subject 7			
Subject 8			
Subject 9			
Subject 10			
Subject 11			
Subject 12			
Conclusion			
Boxouts			
Artwork			
			115

➤ A Greek maths lesson. Sigma is the Greek for "Two beers pronto, señor." I think.

copy must be delivered. I want the spreadsheet to tell me how many days I have to get it done. I've used the cells in the adjacent "F" column for a few workings. I can hide them later to make it look nicer. If you ever type something into a cell and it displays an unexpected result, remember to check the cell formatting.

I've used **DATE2UNIX** to convert both dates into seconds since midnight on 01 January 1970, then subtracted one from the

Subject I	Label printing template	Words	Total
Subject A	Labels have no Linux templates	0	
Subject B	Trial & error blank paper	0	
Subject C	Adjust cell lines to match label edge	0	
Subject D	Save template, copy & paste	191	
Subject E		0	
Subject F		0	
Subject G		0	
Subject H		0	
Subject I		0	
Subject J		0	
Subject K		0	
Subject L		0	
<b>Artwork</b>			<b>191</b>

That Sigma thing working. There are now two layers by my keyboard. Learning languages is always useful.

	A	B	C	D	E	F	G	H	I	J	K
11	<b>Section</b>	<b>Details</b>	<b>Words</b>	<b>Total</b>			<b>Subject B</b>	<b>Details</b>	<b>Words</b>	<b>Total</b>	
12	Introduction		0				Subject A		0		
13	Subject 1		191				Subject B		0		
14	Subject 2		0				Subject C		0		
15	Subject 3		0				Subject D		0		
16	Subject 4		0				Subject E		0		
17	Subject 5		0				Subject F		0		
18	Subject 6		0				Subject G		0		
19	Subject 7		0				Subject H		0		
20	Subject 8		0				Subject I		0		
21	Subject 9		0				Subject J		0		
22	Subject 10		0				Subject K		0		
23	Subject 11		0				Subject L		0		
24	Subject 12		0								
25	Conclusion		0								
26	Boxouts		0								
27	<b>Artwork</b>				191		<b>Artwork</b>				0
28											
29											
30	<b>Subject 1</b>	<b>Label printing template</b>	<b>Words</b>	<b>Total</b>			<b>Subject 9</b>	<b>Details</b>	<b>Words</b>	<b>Total</b>	
31	Subject A	Labels have no Linux templates	0				Subject A		0		
32	Subject B	Trial & error blank paper	0				Subject B		0		
33	Subject C	Adjust cell lines to match label edge	0				Subject C		0		
34	Subject D	Save template, copy & paste	191				Subject D		0		
35	Subject E		0				Subject E		0		
36	Subject F		0				Subject F		0		
37	Subject G		0				Subject G		0		
38	Subject H		0				Subject H		0		
39	Subject I		0				Subject I		0		
40	Subject J		0				Subject J		0		
41	Subject K		0				Subject K		0		
42	Subject L		0				Subject L		0		
43											
44											
45							<b>Artwork</b>				0

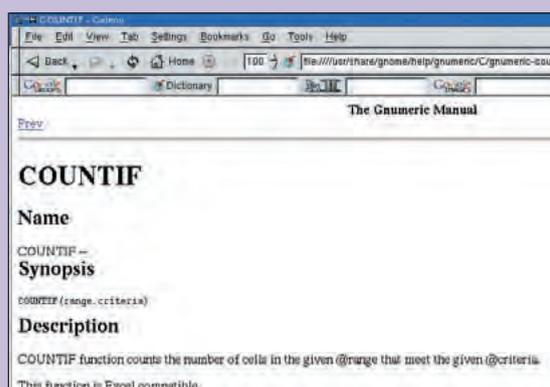
Making all the other parts of the sheet report their values to the summary table.

## Read the manual

Help with HTML anyone?

Gnumeric's manual is a series of HTML pages. There isn't a search function included, but there are a few useful tools to make life easier. There's a command reference that links to related commands. The **COUNTIF** function pictured has hyperlinks to **COUNT** and **SUMIF**.

One advantage with HTML pages is that you can write your own notes onto them if you think they could be improved.



**COUNTIF**

**Name**

COUNTIF -

**Synopsis**

COUNTIF (range, criteria)

**Description**

COUNTIF function counts the number of cells in the given @range that meet the given @criteria. This function is Excel compatible.

The Gnumeric manual.

other and divided by 3600 times 24 (the number of seconds in an hour times the number of hours in a day) to change the result in seconds into days.

I then added a few boxes for the date it was sent, the payment to expect, the date the cheque arrived. To the right of that I put in "Pages needed" and "Word Count" boxes. The Word Count is going to keep track of how the article is progressing. It needs to be at the top, I want all the most important information to be there so that it's easy to see when the file is opened.

Any article is made up of several subject areas. The spreadsheet can help us plan what we intend to write about and ensure that we cover everything. Each subject area will have several points I want to cover. I've selected a suitable area of the sheet and put a border around it (Format>Cells>Border). I'm going to copy and paste this area so I want to make sure that I get all the cells with functions in each time.

## Sigma

Every article will have an introduction, a conclusion, boxouts and artwork. I won't know how many subject areas I will need to cover in an article, but twelve seems a good place to start. I've used the Sigma icon (Sum into the current cell) to highlight a cell where I want all the word counts added together. With that cell highlighted, I just drag a box around the cells to add up the word counts. I've made the area light blue in the screenshot and shown the  $\Sigma$  (Sigma) icon and the Data Entry Area details. I've left a column unused, I might add something later. Note that the "Sum into the current cell" does not include the Artwork line. We will need to return to this when we have put a few more things in.

Select the area inside the border, copy and paste it a few lines below the original. This is going to list our subject details. I've changed the top to read "Subject 1" and made the entries inside read Subject A to L, getting rid of the Introduction and Conclusions. I've made the "details" cell blue and shown the first subject I covered in this article. I probably won't need to include more than twelve points on a subject, but there are a few lines spare.

I can copy and paste this part of the spreadsheet for each of my twelve planned subjects that make up the article. I'll also need one for the introduction, conclusion and artwork. To make it work, I need to make the cells that contain the totals for each subject report the word count to the summary area.

I've used the drawing tool on the left of the work area to tool to draw arrows on the spreadsheet. They start off thin and coloured black. Right-click on them to get a menu to change their properties. The arrows show a couple of cells and how their totals feed into the summary area.

Editors specify how many words they need and it's nice to have something to remind me how many more I need to write. The **IF** function tests a condition and gives alternative outputs if the condition is true or false. I want it to display a message and tell me how many words I need to add or remove. I've asked it to see if cell E7 (the number of words requested) is still greater than cell E8 (the running total of all the work I've done). If it is, I need to write some more copy. The words "Write another" are displayed. I've put the text I want it to display in inverted commas. If I have written too much, it displays "Cut it by". You can put cell references in instead of text and these can be other functions too. The next cell along does a simple subtraction and displays the result. This works fine if I still need to write more copy, but displays a negative value if I have gone over the required number

## Migration

### What's not there if you move from Excel?

*Gnumeric* has a lot of functions and then some. It's a serious maths tool suitable for a wide range of uses. Some things you might miss are *Excel*'s Auditing tool, that shows which cells affect others, the macro, Visual Basic functions and the drawing applet. *Gnumeric* does allow simple shapes and arrows to be included, but you are nearly always better using

a drawing program and importing a picture.

Things you might not miss are the annoying paperclip "help (hindrance)" system, menus that take two clicks to display everything and the price of the software. When you consider that *Gnumeric* is a free download, what you get is an amazingly powerful piece of software.

or words. I've used the ABS function (Absolute value). This makes negative numbers into positive ones and does nothing if they are already positive. The last cell along just has "words" in it to finish the sentence.

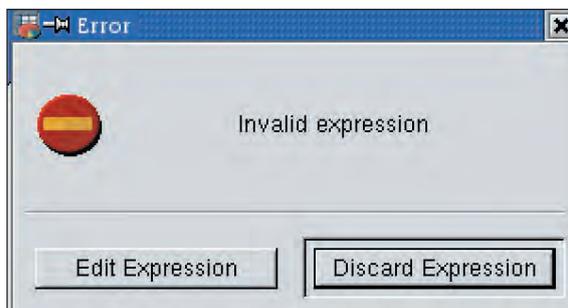
I want to have a box I can tick to show that a cheque has arrived for the article. I've used the Checkbox object from the vertical toolbar to the left of the workspace. Click and drag it where you want it, then right-click to edit its properties. You can put whatever label you want to appear alongside the Checkbox, which returns a value of "TRUE" or "FALSE" to a named cell, depending on if you have clicked the button. I told it to make G7 the named cell, then I made another IF statement below the checkbox. This checks G7 to see if the button has been pressed and displays a message based on the result.

The Function Guru "f(x)" on the toolbar is useful. Highlight a cell with a formula in it and it displays information on how to use it. If you intend to export your spreadsheet for use with *Excel*, you are told if the formula is *Excel* compatible. The **DATE2UNIX** function, as you might imagine, does not have any *Excel* version.

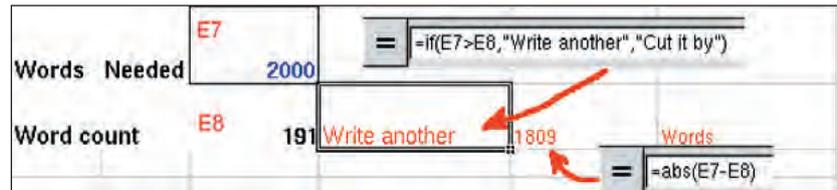
Unfortunately, you have to get the function correct first. If you get something wrong, you get an error message that gives you two options. You can either type it correctly or Discard. The Discard option leaves what you have typed as a text entry, so you can still look up the correct syntax with the Function Guru and put an "=" in to make it a formula when you have got it right.

## Conclusions

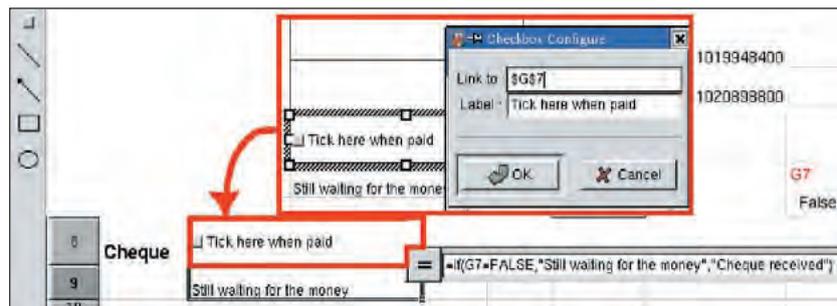
I've shown some of the things that *Gnumeric* can do. The use you get out of it, like any piece of software, depends on your needs and how much effort you are prepared to spend learning what it can do for you. In the next article, we'll do a few graphs and explore a few more useful *Gnumeric* abilities. [LXF](#)



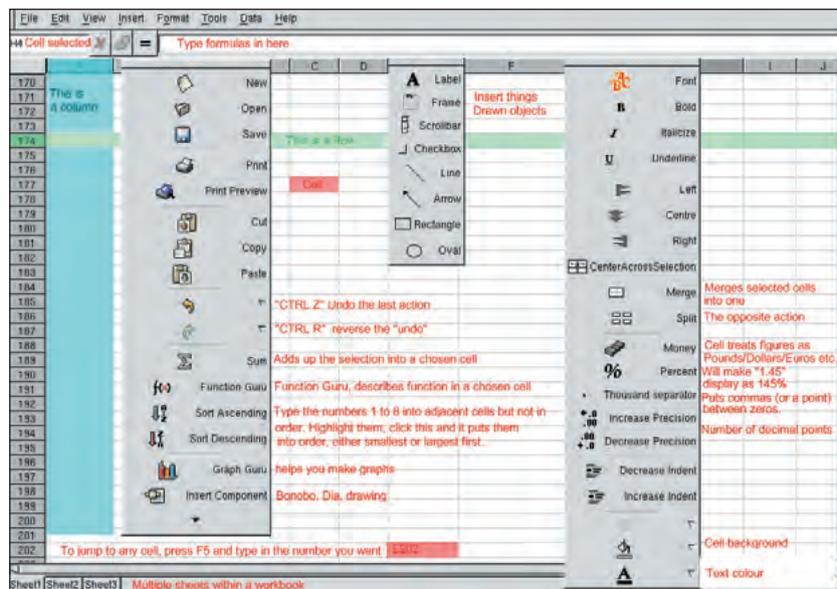
The Function Guru's error message.



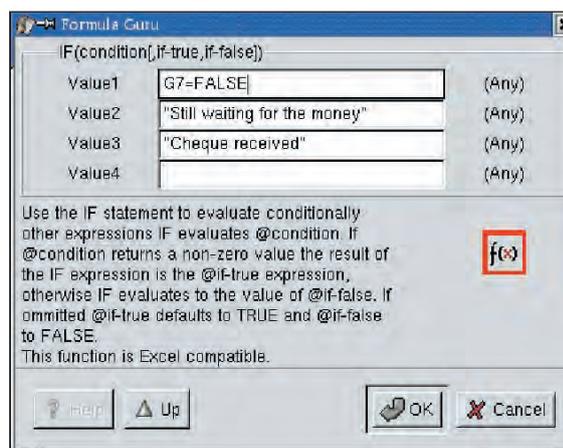
Displaying conditional messages.



The Checkbox can display messages or abuse. The choice is yours. It alters a named cell to either TRUE or "Government Statistics".



Most toolbar icons are familiar, but a few need explanations. Some icons are reserved for future use, they do not actually do anything.



The Function Guru. Like all hippie mystics, it doesn't wash and lives in a world of its own. This is the IF function.

PERL GETS WIDGETS

# Getting graphical



Perl isn't just a text-mangling language. **Charlie Stross** shows how to use the popular Tk graphics toolkit to add graphical front-ends to Perl programs.

**G**UI programming on Unix used to be extraordinarily difficult. As X11 is so low-level, higher level libs were written to provide ready-made widgets (buttons, windows, scroll-bars, and icons). You can bind code to a widget, so that doing something to the widget makes the code run, and you can use "container" widgets (windows or menus) to group child widgets (e.g. scroll bars or menu entries) together.

In the early 90s, Prof John Ousterhout built a graphical toolkit called *Tk*, that provided a very high-level interface to X11 coding. *Tk* doesn't require the programmer to work with X11 input events, monitor the signals from the mouse and re-draw the pointer whenever it moves. *Tk* takes care of the heavy lifting.

These days, *Tk* looks slightly dated – it was designed to look similar to Motif, on UNIX and Linux – but it's a fast, convenient, and easy tool for writing graphical gadgets, and the *Tk* module for Perl lets you write once and run on Windows or \*nix.

## X11 programming basics for Tk

As far as X11 is concerned, the universe consists of windows. A window is a rectangle (or other geometric shape) which contains other windows. The first window, which contains all other windows, is called the *root* window – it's your screen backdrop. An application running under X11 may open other windows, but all of them are contained by the root window.

Windows that contain other windows *manage* their children. A child window is only visible while its coordinates are within the bounds of its parent. Windows, and you, communicate via events, such as keystrokes, mouse button depressions, releases, and movements; a constant stream of X11 events is received by the X server from your input devices, and despatched to the program.

X11 doesn't give you any screen furniture, these features are provided by a Window manager, whose sole job is to interpret the event stream coming from the X11 server and retransmit events to where they're needed. Graphical apps typically start by creating a main window. All the other widgets created by the app live inside it, and receive their controlling events through it. The windows exist in a hierarchy, like a tree, and inherit their attributes from entities closer to the roots.

A widget, such as a button, is actually a whole bunch of windows that work in synchrony. *E.g.*, a button usually consists of six windows – a main enclosing window, four smaller windows that act as bevels, and an inner window. The inner window may have some text rendered on it. Code needs to be tied to the button so that when a mouse-down event is received by the innermost window it (and the bevel windows) changes colour to

reflect the fact that it's in the "down" position, and when a mouse-up is received some other code may be triggered (to do whatever action the button is meant to do). Each of these windows has coordinates associated with each corner, and has to respond correctly when a parent window is dragged (*i.e.* when the parent's X and Y coordinates change, it has to update each child window in the button so that they appear to move together).

## The Tk approach

*Tk* takes the heavy lifting out of this job by providing an object called **button** – a wrapper for all the code that puts a button on the screen. You can send messages to the button object telling it what colour it is, how large it is, what text it carries, what code to execute when it's pressed. Actually, *Tk* provides a lot more than just a button: there's a full set of basic widgets, including windows, menu bars, several different types of button, a canvas (upon which you can draw in various line styles and with various polygons and images), text entry boxes, and the whole panoply of standard GUI elements. (In addition, there are extensions to *Tk* that provide "mega-widgets" – for example, a notebook widget – which consist of a collection of smaller elements configured to work together as a unit, and with their own OO interface.)

Widgets live inside other widgets, just as windows in X11 live inside other windows. Right at the top of the hierarchy in *Tk* is a thing called a **TopLevel** widget – almost invariably a subclass of a fundamental thingy called a window. (Can you see where this is going?) You start writing a *Tk* application by creating the main, top level widget. This acts as a container for scroll bars, windows, buttons, and whatever screen furniture your application needs. These are laid out on screen using a *Geometry Manager* – a set of rules that indicate where widgets go in relation to one another. Finally, you hang subroutines (*callbacks*) off each widget that the user is likely to click on. It's the callbacks that do the back-end heavy lifting associated with your program; *e.g.*, if there's a "Quit" button on your app, and it has files open, the button should have a callback to close the files in good order, then exit the program.

So far we've avoided looking at how the app deals with input. Simple Perl programs typically read text from stdin and write something to stdout. The nature of the input is usually tightly constrained, and it all enters the program at the same place – the line of code that reads stdin. But input to a GUI program can come from a whole range of sources; keyboard, mouse, "resize" events, and so on. And if a user can click the mouse on the "Quit" button, or the "File" menu, or a drawing canvas, how do you handle the possible options?

GUI programs have to respond to input rapidly, from a variety of sources. To cope with this, they are structured around an *event loop*, a case construct. Events come in from the X server. These are matched against windows or widgets that are *registered* with the program, until the one that handles the event is found. Control is handed over to the widget, then returned to the event loop to process the next event.

You may have noticed that widgets are only checked for events if they're *registered*. They're also only visible if they're registered. To register a widget so that the main event loop can draw it on screen and hand events to it, in *Tk* you tell the widget that it belongs to a geometry manager: this acts as the interface between main event loop and widget. Unmanaged widgets aren't visible and hang around, dormant, until a piece of code is called that registers them; then appear (or disappear if we unregister them). This allows us to do things like open up new windows, etc.

## Hello, world!

Here's a "hello.pl" program for you to scratch your head over:

```
#!/usr/local/bin/perl
use Tk;
my $m = new MainWindow;
my $label = $m->Label(
    -text => "Hello, world!"
);
my $button = $m->Button(
    -text => "Quit",
    -command => sub {exit}
);
$label->pack();
$button->pack();
MainLoop;
```

First, we call **use Tk** to load the Perl/*Tk* interface, then create a **MainWindow** object – the top level object in a *Tk* application, which acts as a container for everything else. Here, we call it **\$m**, which is a reference (pointer) to the object. Next, we create an object called a **Label** (a textual tag that goes into a window), and save a reference to it as **\$label**. To create that label inside our **MainWindow**, we call the **Label** method on **\$m**. We're also passing it a parameter, **-text**, with a value of **Hello, World!** – this is the text associated with the label object.

Now we create another object; a **Button**. A **Button** is a widget which, when pressed and released, executes some code. We pass the parameter **-text** to tell our button what its caption is, and the **-command** parameter is a reference to the code to execute when it's released. In this case, we use an anonymous subroutine, but we could define a subroutine elsewhere by name and pass a reference to it: **-command => \&mysub**.

(Remember, we're basically assembling a tree of widgets by hand. They're glued together with references – Perl pointers – and they all hang off the top level widget, **\$m**.)

At this point, we've defined our **MainWindow** and two child widgets, but we haven't registered them with the event loop that dispatches events to them. Nor have we said how large they are, or where to position them in the **MainWindow**. We achieve both these goals by calling **pack()** on each widget. **pack()** is one of the *Tk* geometry managers, which adds the widget to the event



Figure 1: **Welcome to a graphical world! – the result of running our simple "Hello, world!" program.**

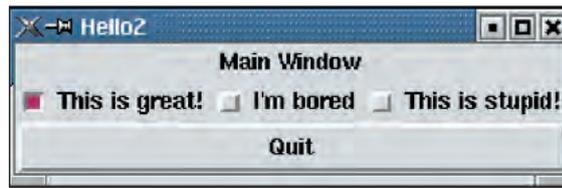


Figure 2: **Buttons fixed in position by the packer.**



Figure 3: **After removing the -fill => x argument...**

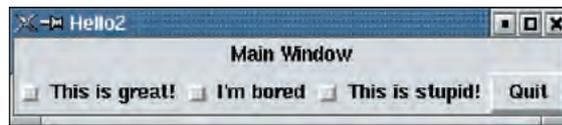


Figure 4: **...and without the before=> \$b1 line.**

loop (so that it can respond to events) and figures out where to draw it on the screen.

Finally: all the preceding stuff is basically setup – declaring widgets, saying what they do when you activate them, registering them, and so on. The *Tk* program is not yet actually displaying anything or responding to events! To set everything in motion, you need to call **MainLoop** – a subroutine that runs the event loop continuously until either the program exits or you call **destroy()** on the **MainWindow** object. (**destroy()** is provided by *Tk* and it causes the event loop to exit and all the widgets to disappear; you can use it if you want to do some non-interactive, non-graphical tidying up while exiting your application.)

What happens if you click on the "Quit" button? The event is passed to the **\$button** object, which executes the callback (subroutine) pointed to by the **-command** parameter. Here the subroutine is a wrapper for **exit**, which causes Perl to exit.

## Geometry managers

Where do your widgets appear on the screen when you register them? Control over positioning is important because without being able to do this, you can't control the layout of widgets in your application's user interface. But you can't make assumptions about the display device; somebody might be using a 640x480 monitor with an 8-bit colour map, while somebody else might have a 1600x1200 display with 32-bit colour.

*Tk* originally provided three geometry managers – ways of managing the positioning of widgets. These are the packer, the placer, and the grid. You specify that a given geometry manager has control of a widget by calling it as a method on the widget:

```
my $cb = $m->CheckButton(-text => "red");
$cb->pack()
```

This creates a checkbutton widget with the text **red**, and tells it that it's controlled by the packer (without supplying any options to tell the packer how to behave).

Each geometry manager displays its widgets in a different way. The packer assigns invisible bounding boxes to a widget; these are anchored (either to other widgets under packer control, or to the toplevel window), and can be resized. The widgets show up within non-overlapping boundaries defined by the boxes. The grid defines a row/column table and lets you place widgets in specific positions in the grid. (This is great for spreadsheet-like layouts.)

Unlike the other two managers, which do not permit widgets



◀◀ to overlap and which define widget positions relative to some other entity (allocation boxes in the case of the packer, the table layout in the case of the grid), the placer allows you to place widgets at absolute coordinates within the parent window. It allows buttons and other widgets to overlap, and enables all sorts of confusion – but it's the most appropriate manager to use for generating graphics on the fly, rather than simple user interfaces that rely on widgets at fixed positions.

Here's a simple test program that creates a bunch of buttons, and that demonstrates how we use the packer (see **fig 2**):

```
#!/usr/local/bin/perl
use Tk;
my $m = new MainWindow;
my $label = $m->Label(-text => "Main Window")->pack;
my $b1 = $m->Checkbox(
    -text => "This is great!",
    -command => sub { print STDERR "This is great!\n"; }
);
my $b2 = $m->Checkbox(
    -text => "I'm bored",
    -command => sub { print STDERR "I'm bored!\n"; }
);
my $b3 = $m->Checkbox(
    -text => "This is stupid!",
    -command => sub { print STDERR "This is stupid!\n"; }
);
my $exit = $m->Button(-text => "Quit",
    -command => sub { exit }
);
$b1->pack( -side => "left",
    -expand => 1
);
$b2->pack( -side => "left",
    -expand => 1
);
$b3->pack( -side => "left",
    -expand => 1
);
$exit->pack( -side => "bottom",
    -expand => 1,
    -fill => "x",
    -before => $b1
);
MainLoop;
exit;
```

There are several points to note about this program before we examine how it uses the packer. First, we can append the **pack()** commands to each widget when we create the thing – the reason we don't do so in this example is to make it easier to read. Each of the **Checkbuttons \$b1 to \$b3** has a simple callback that makes it print some text on the standard output when we check it; this is a handy debugging technique you can use. The **Checkbuttons** are not grouped together and they don't keep track of any underlying state – but you could use the **-command** callbacks to make them update an underlying data structure and, if necessary, spawn a new widget that would show up as soon as it is managed by one of the geometry managers.

**pack()** allocates space in the toplevel widget from the remaining area which is not already occupied by widgets. By calling **pack()** with the **-side => left** argument, we are telling it that the current widget is to be packed to the left of the available

space. Subsequent widgets are packed into the window in whatever space is left, and the **-expand => 1** parameter tells the packer to allow the widgets to expand to fill the available space.

The **\$exit** widget is treated differently. First, we tell the packer to anchor it to the bottom of the window. The toplevel window expands to the minimum size necessary to accommodate all the managed widgets in the order in which they are packed. Because the first three items are packed to the left, in a row, the **\$exit** widget ought to be positioned against the bottom of the window – and to the right of the row of checkbuttons. But we also specified **-before => \$b1**. This tells the packer to put **\$exit** first in its list of widgets to allocate space for in the window. And we gave the **-fill => x** argument; this tells the packer to allow the button to expand to fill the X-axis of its allocation rectangle.

If we delete the **-fill => 1** line, we get **fig 3**, and if we delete the **before => \$b1** line, we see **fig 4**.

When you take an empty window and tell the packer to allocate space for a widget at the left, it creates an invisible allocation rectangle that occupies the entire vertical height of the window, at the left hand side of the window. If you tell it to pack against the top, the allocation rectangle is as wide as the window and occupies the top of it. Here's a small program (from *Learning Perl/Tk*) that demonstrates where buttons end up if we pack them to left, right, top and bottom (see **fig 5**):

```
#!/usr/local/bin/perl
use Tk;
my $m = new MainWindow;
my $label = $m->Label(-text => "Main Window")->pack;
my $b1 = $m->Button(
    -text => "Top",
    -command => sub { exit; }
);
my $b2 = $m->Button(
    -text => "Left",
    -command => sub { exit; }
);
my $b3 = $m->Button(
    -text => "Right",
    -command => sub { exit; }
);
my $b4 = $m->Button(
    -text => "Bottom",
    -command => sub { exit; }
);
$b1->pack( -side => "top",
);
$b4->pack( -side => "bottom",
);
$b2->pack( -side => "left",
);
$b3->pack( -side => "right",
);
MainLoop;
exit;
```

It's pretty ugly because the allocation rectangles aren't filled. If we tell *Tk* to allow widgets to expand to fill their allocation rectangles (albeit without overlapping), like this (see **fig 6**):

```
$b1->pack( -side => "top",
    -fill => "both"
);
```

Note that the reason we get the big buttons at top and



Figure 5:  
Packing the buttons to the left, right, top and bottom – a demonstration of **pack()**.



Figure 6:  
We get a much prettier result if we allow the widgets to fill their allocation rectangles.

bottom is that we packed them – added them to the packer's list of managed widgets – before we packed the left and right widgets. Try experimenting by changing the order in which the widgets are packed (just cut and paste the `pack()` commands into a different order). You may find the results a little surprising!

The packer doesn't just position widgets in a window relative to one another; it can be used to destroy widgets:

```
$b3->packForget()
```

This causes widget `$b3` to vanish from the packing order and the window. The packer can return info on a widget using `packInfo()`:

```
@info = $b3->packInfo();
```

```
print "[", join("]", @info), "]"
```

Which returns something like this:

```
[-in][MainWindow=HASH(0x100f0bb8)][-anchor][center]
```

```
[-expand][0][fill][both][ipadx][0][ipady][0][padx][0]
```

```
[-pady][0][side][right]
```

The info here includes a reference to the parent widget, as well as all the applied attributes of the widget we're querying.

We can also use `packSlaves` to retrieve a list of references to all the child widgets managed by a parent that's using the packer:

```
my @info = $m->packSlaves();
```

```
print "[", join("]", @info), "]"
```

Which produces something like this:

```
[Tk::Label=HASH(0x102a91b0)]
```

```
[Tk::Button=HASH(0x102c36f0)]
```

```
[Tk::Button=HASH(0x102c38c4)]
```

```
[Tk::Button=HASH(0x102c36a8)]
```

```
[Tk::Button=HASH(0x102c384c)]
```

Using these returned references we can allow a parent to tweak its child widgets attributes – e.g., by using the `ipadx` and `ipady` methods to re-pack them in a larger bounding box.

## Going Further

This tutorial barely scratches the surface of what you can do with Perl/Tk. We'll look at some more widgets in the next tutorial. One topic we haven't even looked at is how to go about designing a GUI that does something useful. Here's a final example which acts as a front-end to the `wc(1)` word count program. When you run it, you get two buttons: a "pick a file" button and a "quit" button (see **fig 7**). "Pick a file" leads you to select a file (**fig 8**). After picking a file, the "pick" button vanishes and is replaced by the number of characters, words, and lines in the file (**fig 9**).

This program shows a number of useful features – using the `geometry()` method to tell the toplevel window where to position itself, using `Tk::FileSelect` (external mega-widget that provides a file selection box) to grab a file, and destroying a widget and replacing it with another on the fly using `packForget()`. The callback does the heavy lifting, while the GUI is almost entirely independent of the code that actually does stuff.

```
#!/usr/local/bin/perl
use Tk;
use Tk::FileSelect;
$rootdir = shift @ARGV || `pwd`;
chomp $rootdir;
$name = "";
$m = new MainWindow;
$label = $m->Label(-text => "Word Count");
$pickbutton = $m->Button(-text => "Pick a file",
    -command => sub{$label = &getfile($m, $label)}
);
$quitbutton = $m->Button(-text => "Quit",
```

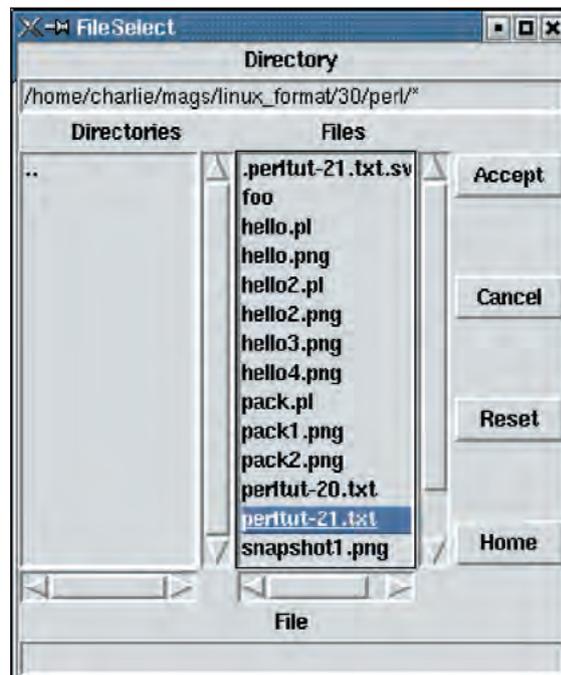


Figure 8: Selecting a file through our frontend.

```
-command => sub {exit}
);
$m->geometry("+0+0");
$label->pack(-side => "top",
    -fill => "both"
);
$pickbutton->pack(-side => "top",
    -fill => "both"
);
$quitbutton->pack(-side => "top",
    -fill => "both"
);
MainLoop;
exit;
sub getfile {
    my ($m) = shift @ ; # the toplevel widget
    my ($label) = shift @_ if @ ; # the label widget
    if ($label) {
        $label->DESTROY;
    }
    my ($fref) = $m->FileSelect(-directory => $rootdir);
    my ($filename) = $fref->Show;
    my $swc = join("\n", `wc $filename`);
    chomp $swc;
    my $label = $m->Label(-text => "wc stats: $swc");
    $label->pack(-fill => "both",
        -expand => 1,
        -side => "top",
        -before => $pickbutton
    );
    $pickbutton->packForget();
    return $label;
}; LXF
```

Figure 9: The results from our graphical frontend to the word count program.



## REGULAR EXPRESSIONS

# Speaking Java

Make textual search and replace operations in Java easy with regular expressions. **Richard Drummond** tells all.



Last time we looked at how to parse a set of records stored as an XML file and build a corresponding set of Java objects in memory to hold the data contained in those records. If you remember, this is used for the back-end for our mini project: we are constructing a Java applet which can be used to search an index of files on a Linux Format coverdisc. The code we have implemented so far scans an index file and marshals the data it contains as a list of `IndexedPackage` objects (see the box, right). Now that we are able to load the data, we need to think about how we are going to search it.

## Finding what we want

What we want to do is to be able to present the user with a GUI with which they can enter a query and then display as a result the packages on disk which match that query.

We could imagine that different types of query might be needed. In the simplest case, a simple keyword query – which lets the user search for a given keyword in package names and package descriptions – should be possible. The GUI for this type of query could be presented to the user as a single string gadget, and hence would be compact and simple to use – just the thing for casual searches. However, more complex queries should also be allowed. For example you might want to list all the packages whose names contain the substring 'KDE', which are on discs 28 or 29 and which have a version number greater than version 3.0. Of course, this type of thing will require a much more complex user interface.

Leaving aside the user interface for just now, it should be obvious to you that we can implement the actual back-end code to perform a query just once and use it for both types of query: the simple query is just a restriction of the more complex query type. If our search engine is flexible enough, it can do both.

In the case of the simple keyword search, what is actually being performed is a search of the more general type where we are looking for any substring in the `Name` and `Descriptions` fields of each `IndexedPackage` object that matches the given keyword.

With that brief architectural overview out of the way, it's time to think about the mechanics of searching. While we could no doubt implement our own code to match an `IndexedPackage` object against various criteria, we are far too lazy to do this. And,

**Our search engine should be able to support queries of differing complexities.**

## The IndexedPackage class

A reminder just in case you've forgotten

The `IndexedPackage` class represents a single software package in an index. This extract will refresh your memory of its fields.

```
class IndexedPackage
```

```
{
    String Name;      /* The package name */
    String Disc;     /* Which disc was it on? */
    String Description;
    String Path;     /* Location on disc */
    String Version;  /* Version number */
    String Link;     /* Website URL */
```

```
/* implementaton */
```

```
}
```

anyway, object-oriented programming is about code re-use, after all. For maximum flexibility and minimum work on our part, we are going to base our searching code on regular expressions and use an existing regular expression toolkit.

## Regular expressions

A regular expression is a wild card pattern which matches some set of strings. Regular expressions are used in many text-processing utilities – such as *grep*, *awk* and *sed* – and hence many syntaxes for specifying regular expressions have evolved over the years. Popular schemes include Perl, *sed* and *awk* regular expressions.

With the release of JDK 1.4, Java now has regular expression support as standard, implemented in the package `java.util.regex`. While it's nice to have a standard, there are many other regular expression toolkits for Java (see *Other regex implementations*). These have their own particular advantages and disadvantages, and a third-party toolkit is especially useful when you can't rely on your code being deployed on JDK 1.4. For just now, however, we will concentrate on JDK 1.4's regular expressions only.

For simple, one-off patterning matching, JDK 1.4 provides some additional methods in the `String` class which make using regular expressions in Java extremely easy. The `matches()` method accepts a regular expression as an argument and returns true if the string matches that pattern and false otherwise. The syntax for specifying regular expressions in JDK 1.4 is very like Perl's regular expression syntax. It is worth noting, however, that in the `String.matches()` method, the pattern must match the whole of the string, not simply any substring. So, for example, if we declare the string

```
String s = "Linux Format is great"
```

then the expression

```
s.matches("Format")
```

will evaluate as false, but the expressions

```
s.matches( ".Format.*" ) // matches any string containing
                        // the substring "Format"
```

```
s.matches( "(\\w+\\W+){3}\\w+" ) // matches any string with four
                                // words but which does not
                                // begin or end on whitespace
```

will both evaluate as true. If you are not familiar with the syntax of Perl's regular expressions then the box *Java regex syntax* lists some of the most common. Note that Java, like Perl, treats a backslash in a string literal as an escape character: to actually include a backslash in a string, you must use double backslashes.

Also useful in JDK 1.4 are the methods `String.replaceNext()` and `String.replaceAll()` – which replace the next occurrence or all occurrences, respectively, of a given pattern in a `String` with a specified substring – and `split()` – which returns an array of substrings of a string, separated by delimiters that match the specified pattern.

As an example of these, if we declare the strings `s` as follows

```
String s = "Linux and java are a great combination"
```

then the expression

```
s.replaceAll( "\\bjava\\b", "Java" )
```

will replace all occurrences of the word 'java' with word 'Java'.

The expression

```
s.split( "[ \\t\\n\\r]+" )
```

will split the string `s` into words, using whitespace as a delimiter (escape codes are specified as normal in a regular expression) and thus evaluates as an array containing the strings 'Linux', 'and', 'java', 'are', 'a', 'great' and 'combination'.

## Java Regex syntax

### How to specify regex patterns in JDK 1.4

The following table lists some of the constructs available in JDK 1.4's regular expression matching. For a full list, consult the API manual. `java.util.regex`'s pattern matching is very similar to Perl 5, so any tutorial on Perl regular expressions should also prove helpful.

Construct	Matches
<code>x</code>	The character <code>x</code>
<code>.</code>	Any character
<code>[abc]</code>	The characters <code>a</code> , <code>b</code> , or <code>c</code>
<code>[^def]</code>	Any character except <code>d</code> , <code>e</code> , or <code>f</code>
<code>[a-z]</code>	The range characters from <code>a</code> to <code>z</code> (inclusive)
<code>\\d</code>	A digit
<code>\\D</code>	Any non-digit
<code>\\s</code>	A whitespace character
<code>\\S</code>	A non-whitespace character
<code>\\w</code>	A word character
<code>\\W</code>	A non-word character
<code>^</code>	The beginning of a line
<code>\$</code>	The end of a line
<code>\\b</code>	A word boundary
<code>\\B</code>	A non-word boundary
<code>X?</code>	<code>X</code> , once or not at all
<code>X*</code>	<code>X</code> , zero or more times
<code>X+</code>	<code>X</code> , one or more times
<code>X{n}</code>	<code>X</code> , exactly <code>n</code> times
<code>X{n,}</code>	<code>X</code> , at least <code>n</code> times
<code>X{n,m}</code>	<code>X</code> , at least <code>n</code> times but not more than <code>m</code> times
<code>XY</code>	<code>X</code> followed by <code>Y</code>
<code>X Y</code>	Either <code>X</code> or <code>Y</code>

## Other regex implementations

If you don't have JDK 1.4, check out these other regex libraries

### Jakarta Regexp

<http://jakarta.apache.org/regexp/>

Part of the Jakarta project, this is a simple but mature regular expression library.

### Jakarta ORO

<http://jakarta.apache.org/oro/>

The Jakarta-ORO Java classes are a set of text-processing Java classes that provide Perl5 compatible regular expressions, AWK-like regular expressions, glob expressions, and utility classes for performing substitutions, splits, filtering filenames, etc.

### gnu.regex

<http://www.cacas.org/java/gnu/regexp/>

A full-featured regular expression package for Java distributed under the LGPL. It supports most features of Perl5 regular expressions, as well as support for many other regular expression syntaxes, including awk and Emacs.

### JRegex

<http://jregex.sourceforge.net/>

A compact and efficient regular expression library for Java distributed under the BSD license. It provides support for Perl5.6 regular expression syntax and has full Unicode support.

## Re-using patterns

The `java.util.regex` package contains classes which are more suited for repeated application of a pattern expressed as a regular expression.

The `Pattern` class forms the basis of `java.util.regex` and encapsulates a regular expression. Its class method `compile()` accepts a `String` containing a regular expression as an argument, compiles that expression into a more efficient form ready for use, and returns an instantiated `Pattern` object for you to use. You can then invoke that object's `matcher()` method to apply the pattern to a specific sequence of characters stored in a `String`, `StringBuffer`, or `CharBuffer`.

The `Matcher` object that is returned by `Pattern.matcher()` implements the methods `matches()`, `replaceFirst()` and `replaceAll()` just like the `String` class does, and can be used similarly to test whether the pattern matches or to perform search and replace operations.

The `Matcher` class has in addition many more useful tools, though. The `find()` method, for example, can be used to locate the next subsequence of characters in the given sequence which match the specified pattern. This method returns true on successfully finding a match – in which case the `group()` and `start()` methods can be used to access the substring which matched the pattern and the position of that substring in the character sequence being searched, respectively. Have a look at this extract.

```
String s = "Linux Format is the UK's best-selling news-stand
magazine dedicated to Linux.;"
```

```
Pattern p = Pattern.compile( "\\bLinux\\b" );
```

```
Matcher m = p.matcher( s );
```

```
while( m.find() )
```

```
System.out.println( "Found ""+m.group()+"" at position
"+m.start() );
```

We have just had enough space here to give a brief introduction to JDK1.4's regular expression matching, but from what little we've covered you should begin to see how powerful it is. Find out more by visiting the Java documentation page at <http://java.sun.com/j2se/1.4/docs/>. Next time we will look in detail at how we can use regular expressions in our index browser applet. And, there, because we don't want to have to rely on the user having JDK 1.4 installed, we shall be using one of the other regular expression libraries for Java. Don't worry, though. They all work rather similarly. [LXF](#)



## OOP SYNTAX

# OOP syntax in Object Pascal

**PART 12:** This month **Brian Long** looks at the Object Pascal support for building your own objects in Kylix applications.

Last month we had some fun doing some simple animation. In the example provided we had a single layer of animation drawn across a background. It is very possible to use additional bitmaps to build extra animation layers, just as with cartoon animation.

We might come back to this area later in the series and see what we can achieve, but this month we focus our attention on OOP (Object Oriented Programming) support in Object Pascal.

## OOP fundamentals

Object oriented programming is really a means to allow programmers to mirror the real world a little more closely than may be possible otherwise. If you have the ability to create a model in your program that closely mirrors what you are trying to represent then the theory goes that your job will be easier and the resultant code will be easier to maintain.

*Kylix* is an OOP tool, which means it allows you to write classes that exhibit the three so-called cornerstones of OOP: encapsulation, inheritance and polymorphism. We'll try and get an understanding of each of these principles this month whilst developing some sample classes. Additionally we'll see how properties are defined in order to provide "storage-with-side-effects."

To start off we should get a few basics out of the way. An object is a programmatic data entity that contains everything required to represent something. What the "something" is completely dependent on the application and the object, but is often something from the real world that is being modelled to a greater or lesser extent.

In the real world any arbitrary object (for example, a pencil) has various attributes (such as the length, mass, density, lead hardness and bluntness) and behaviour (such as the ability to

make a mark on something when the tip makes contact). To represent a pencil in a program you first define a *class* type.

A class is a data type that defines the template that all pencils should follow. It defines the attributes as data fields and/or properties and the behaviour as methods. Once the class has been defined you can create instances of the class; each instance is called an *object*.

A component in *Kylix* is an object that has built-in behaviour to allow it to interact with the *Kylix* IDE (notably the **Form Designer** and **Object Inspector**). Components can be stored in form files when you save your project, which are compiled into your executable during the compile/link cycle. This allows the form that you set up at design-time to be recreated as needed when the program is running.

Nevertheless, components are objects and so follow the definition of an object. For example, a **TEdit** component represents an edit widget. It has properties to reflect the widget attributes such as **Text**, **Top**, **Left** and **Color**. It also has methods to surface the widget behaviour such as **Clear**, **PasteFromClipboard** and **Hide**.

When you design a form using the **Form Designer**, *Kylix* defines a form class to allow programmatic access to the form's attributes (including all the components you drop on it) and behaviour (the event handlers you set up for the form and the components on it). A form is actually a special type of component.

*Kylix* comes with many classes that do not meet the criteria to be components or forms and so are not component classes; they are simple object classes. Since component classes add in extra complications we will concentrate on simple object classes in this foray.

## Encapsulation

Let's start our look at class-writing syntax by defining a simple class that represents a car. *Encapsulation* allows us to define the

data attributes and any related routines that typically act upon that data in the same class, rather than declaring global variables and global routines.

The basic type definition looks like this and would be found in the interface section of a unit (unless you didn't want other units to access it, in which case it would be in the implementation section):

```
type
  TCar = class
end;
```

This defines a basic class called **TCar**, but which is empty. Note the convention of prefixing any type definitions with **T** (to imply it's a type). We'll add data attributes to represent the speed of the car and whether it's parked and a couple of methods that let the car be parked and driven. The method implementations must be placed in the implementation section of the defining unit, regardless of which section the class is defined in.

```
type
  TCar = class
    Speed: Integer;
    Parked: Boolean;
    procedure Park(ParkCar: Boolean);
    procedure Drive;
  end;
...
uses
  QDialogs;

procedure CarError(const ErrorMsg: String);
begin
  ShowMessage(ErrorMsg)
end;

procedure TCar.Drive;
begin
  ShowMessage('Car is driving. Speed is %d', [Speed])
end;

procedure TCar.Park(ParkCar: Boolean);
begin
  if ParkCar <> Parked then
    if ParkCar and (Speed > 0) then
      CarError('Car must be stationary to be parked')
    else
      Parked := ParkCar
  end;
end;
```

This is a simple example and there are various potential integrity problems with it.

## Data hiding

It is generally accepted that when writing classes you should employ *data hiding* to protect the integrity of data items. Take the class above; there is nothing to stop any code in an application changing the value of the **Speed** field to a ridiculously high value. Also, there is nothing to stop some code directly changing **Parked** to **True**, even if **Speed** is set to **100**.

In other words, when data fields are directly exposed from a class, there is no protection against their values being modified in error. To avoid this kind of error requires a mechanism to hide the data items from code that uses objects of your class type.

Instead, code outside the class must use routines to access the data. There is a routine to read the data value (the *getter* routine) and another to update the value (the *setter* routine).

The setter can employ validation code to ensure no invalid value is assigned to the underlying field. The getter typically just returns the underlying field value, but can run any additional code you feel may be necessary when returning the value.

## Data access control

There are various levels of access restriction you can assign to your class members. Of interest right now are the private and public access levels. There are two other levels, but we will look at those later.

Defining something as public allows code outside the class to access it, whereas marking it private prevents any code outside the unit that defines the class from accessing it. Following this idea we might modify the class to look like this:

```
type
  TCar = class
    private
      FSpeed: Integer;
      FParked: Boolean;
    public
      procedure SetParked(Value: Boolean);
      function GetParked: Boolean;
      procedure SetSpeed(Value: Integer);
      function GetSpeed: Integer;
      procedure Drive;
    end;
...
uses
  QDialogs;

procedure TCar.Drive;
begin
  ShowMessage('Car is driving. Speed is %d', [FSpeed])
end;

function TCar.GetParked: Boolean;
begin
  Result := FParked
end;

function TCar.GetSpeed: Integer;
begin
  Result := FSpeed
end;

procedure TCar.SetParked(Value: Boolean);
begin
  if Value <> FParked then
    if Value and (FSpeed > 0) then
      CarError('Car must be stationary to be parked')
    else
      FParked := Value
  end;
end;

procedure TCar.SetSpeed(Value: Integer);
begin
  FSpeed := Value
end;
```

◀◀ Note that a convention is being applied here where private data fields start with an **F** (for field).

With this implementation no outside code can change the **FParked** value without going through the **SetParked** routine. Similarly nothing can change the car speed without going through **SetSpeed**. Admittedly, **SetSpeed** does no validation (such as limiting its top speed) but we'll add that in soon. The important thing is that code has to execute in the class to change the value of a class field.

## Properties

One downside to this data hiding approach is that you have two different routines to call to read and write a data member of a class (such as **SetSpeed** and **GetSpeed**). Object Pascal simplifies this issue by introducing support for properties.

A property is a mechanism that exposes an item from a class that looks just like a data attribute, but which is defined in terms of what happens when it is read from or written to. You can define a property to call a getter and setter routine when it is accessed, as in:

```
property Parked: Boolean read GetParked write SetParked;
property Speed: Integer read GetSpeed write SetSpeed;
```

This would allow you write code like this:

```
if not Car.Parked then
```

```
begin
```

```
  if Car.Speed <> 0 then
```

```
    Car.Speed := 0;
```

```
    Car.Parked := True
```

```
end;
```

instead of this:

```
if not Car.GetParked then
```

```
begin
```

```
  if Car.GetSpeed <> 0 then
```

```
    Car.SetSpeed(0);
```

```
    Car.SetParked(True)
```

```
end
```

You can read and write using a single identifier (the property) rather than calling two routines (the getter and setter).

However, if the getter would simply return a private data field, you can specify that the property returns it directly and save the trouble of a pointless routine. Similarly if you need no validation when setting the value, you

can specify that the property writes directly to the data field.

You can also define read-only and write-only properties and there is also support for array properties (implemented with the array indexer passed to the getter/setter as an extra parameter) and indexed properties (where several properties are implemented with the same getter/setter, which is passed an index value to identify the property).

We can rewrite the class like this (the method implementations that remain are exactly the same as before):

```
type
  TCar = class
```

```
  private
    FSpeed: Integer;
    FParked: Boolean;
  procedure SetParked(Value: Boolean);
  procedure SetSpeed(Value: Integer);
  public
    procedure Drive;
  property Parked: Boolean read FParked write SetParked;
  property Speed: Integer read FSpeed write SetSpeed;
end;
```

## Inheritance

Our car class may well do a simple job of representing a car, but things are more interesting and slightly more realistic if we introduce some classes that represent particular types of car. We can simplify this task by using *inheritance*.

When you hear someone being passionate about OOP you can be sure the word “reuse” will crop up at some point. Reuse is enabled by inheritance which allows a new class to be defined as a superset of an existing class. If you can locate a class that provides a proportion of what you need from a new class, you can make your new class inherit the existing class, which represents your new class's starting point. The old class is called the *ancestor* of the new class; the new class is a *descendant* of its ancestor.

In a rich class library (such as *CLX*) there will be many classes that inherit from other classes. When designing a class hierarchy for an application, careful thought should be put to what classes exists and what inherits from what. Any common code should be implemented in a class, and the various specifics can be implemented in descendants of the common ancestor. For example, all components inherit from the **TComponent** class, which defines all the functionality and properties necessary to be worked with at design-time.

All Object Pascal classes you define must inherit from **TObject**, either directly or by inheriting something that is already inherited from **TObject**.

**TObject** is a base class of any Object Pascal class hierarchy – it is the ultimate ancestor of every class and so sits right at the top of the hierarchy (or at the very bottom, depending which way up you picture it). It provides various simple services required by the *Kylx* RTL (**BaseCLX**). For example every class has a **ClassName** method inherited from **TObject** that returns the class name as a string. So, for example, **Button1.ClassName** returns the string **TButton**.

**Figure 1** shows a representative illustration of the inheritance that leads to the **TBitBtn** class.

The classes we have looked at so far have not specified an ancestor and so implicitly inherit directly from **TObject**. You can make this explicit by changing the class definition to:

```
type
  TCar = class(TObject)
  private
    ...
  end;
```

Note that Object Pascal implements *single inheritance*, which means a class can inherit directly from a single class, not from multiple classes, as you can in C++, which supports *multiple inheritance*.

We can build a small hierarchy that allows specific cars to be implemented using **TCar** as a base class and inheriting its generic

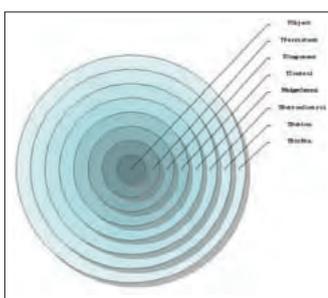


Figure 1: A representation of the inheritance present in a **TBitBtn**.

car behaviour. Each descendant will introduce the specifics of that type of car.

Additionally we can, if we wish, define custom exception classes inherited from `Exception`. These exceptions can be raised to represent any error conditions. So we can re-implement `CarError` and `TCar.SetParked` as:

```
type
  TCar = class(TObject)
  ...
end;

ECarError = class(Exception);
...
procedure CarError(const ErrorMessage: String);
begin
  raise ECarError.Create(ErrorMessage)
end;

procedure TCar.SetParked(Value: Boolean);
begin
  if Value <> FParked then
  begin
    if Value and (FSpeed > 0) then
      CarError('Car must be stationary to be parked');
    FParked := Value
  end
end;
```

## Data access control, continued . . .

With an understanding of inheritance we can now properly cover the data access levels. These are defined in terms of how they affect a class writer and an object user. A class writer is at liberty to inherit from any class you define, and an object user can create an instance of your class and try to access its members.

You declare something as *private* if you wish to exclude it from class writers and object users, meaning it is only for your use. Data fields are commonly declared private.

You declare something as *protected* if you wish to exclude it from object users, but still leave it available for class writers. Property getters and setters whose implementations need to be extended or modified by descendant classes will typically be declared protected.

You declare something as *public* if you wish to allow class writers and object users to have access to it. Anything you want accessible at runtime is typically declared public.

You declare a component property as *published* if you wish to exhibit the same behaviour at runtime as if it were public, but wish it to appear on the **Object Inspector** at design time. Since we are not writing components here, we have no need for this access level.

Note that to be perfectly accurate, private and protected items can be accessed from other code in the same unit as the class, however it is typical that the class writer's class or object user's code would be in a different unit from your class.

## Constructors And Destructors

Before implementing the descendants let's extend the base class a bit more. Since any given car has a maximum attainable speed we should make this information available (but not modifiable) and also adhere to it. Also the car should start its life in a parked state.

## Looking Towards Kylix 3

### C++ and Object Pascal

Borland's annual conference (BorCon 2002) was held this year in May in Anaheim, California. Of interest to *Kylix* developers (and also developers who like C++, but don't use *Kylix*) was information about the next version of *Kylix* that was given in the Product Address session.

It seems that *Kylix 3* (codenamed *Cortez*) will include the Object Pascal language support that has been present in *Kylix 1* and *Kylix 2*, but will also include the first version of Borland's C++ compiler for the Linux platform. This means that

*Kylix* developers will have the choice of writing *CLX* applications either in Object Pascal or C++.

Additionally, Linux programmers who are familiar with the Linux and/or X windows APIs can write non-*CLX* applications as they have done previously, but in a fully functional integrated development environment. *Kylix 3* is expected to be released later this year, so we'll keep you up to date with news about it. At this stage there has been no information on what will be offered in the *Open Edition* of *Kylix 3*, but we'll keep our ear to the ground.

We can achieve this by setting the top speed and the parked value in the constructor, which is a special method called to initialise a freshly created object. Every *Kylix* class has a constructor called **Create**; the one inherited from **TObject** takes no arguments but many classes redefine the constructor to accept appropriate arguments to initialise the object. Additionally, a read-only property can make the top speed available.

Notice in the code below the constructor calls the inherited constructor to ensure that any initialisation code inherited from the ancestor is not lost.

Also note that the **SetTopSpeed** protected method is there to allow class descendants to still get a value assigned to **FTopSpeed**. If a **TCar** descendant were defined in some other unit, it would not have access to the **FTopSpeed** private data field.

```
type
  TCar = class(TObject)
  private
    FTopSpeed: Integer;
  ...
  protected
    procedure SetTopSpeed(Value: Integer);
  public
    constructor Create;
    property TopSpeed: Integer read FTopSpeed;
  ...
end;
...
constructor TCar.Create;
begin
  inherited Create;
  //Descendant car classes will set their top speeds
  SetTopSpeed(0);
  Parked := True;
end;

procedure TCar.SetSpeed(Value: Integer);
begin
  if Value <> FSpeed then
  begin
    if Value > TopSpeed then
      CarError(Format('Top speed of %d cannot be exceeded',
        [TopSpeed]));
    if Parked and (Value <> 0) then
      CarError('Cannot drive a parked car');
    FSpeed := Value
```



```

end;
end;

procedure TCar.SetTopSpeed(Value: Integer);
begin
  FTopSpeed := Value;
end;

```

When the object is destroyed by a call to its **Free** method another special method, the destructor, is called to do any tidying up needed. Every *Kylux* class has a destructor called **Destroy**.

Constructors and destructors are examples of methods that should definitely call their inherited versions to ensure important setup/tidy-up code is still executed. For other methods, whether you call the inherited version depends on the circumstances. If you do call the inherited code, you will be extending the behaviour of the ancestor's method, otherwise you will be changing it.

## The Classes

We can now implement the descendant classes, which will set the top speed in their constructor and implement custom driving behaviour.

```

type
  TPorsche911 = class(TCar)
  public
    constructor Create;
    procedure Drive;
  end;

  TMorrisMinor = class(TCar)
  public
    constructor Create;
    procedure Drive;
  end;
...
constructor TPorsche911.Create;
begin
  inherited Create;
  SetTopSpeed(174); //top speed of some model of 911
end;

procedure TPorsche911.Drive;
begin
  ShowMessage('Driving in a Porsche 911 at %d mph (vroom, vroom!); [Speed]
end;

constructor TMorrisMinor.Create;
begin
  inherited Create;
  SetTopSpeed(62); //top speed of some model of the MM
end;

procedure TMorrisMinor.Drive;
begin
  ShowMessage('Driving sedately in a Morris Minor at %d mph',
[Speed])
end;

```

With these classes complete we can test them. A test application is supplied on this month's disk called *ObjectsEg.dpr* that has the following event handlers set up for the form's

**OnCreate** and **OnDestroy** events and the **OnClick** events of a couple of buttons:

```

procedure TForm1.FormCreate(Sender: TObject);
begin
  Randomize;
  Morris := TMorrisMinor.Create;
  Morris.Parked := False;
  Porsche := TPorsche911.Create;
  Porsche.Parked := False;
end;

procedure TForm1.FormDestroy(Sender: TObject);
begin
  FreeAndNil(Morris);
  FreeAndNil(Porsche);
end;

procedure TForm1.btnDriveMorrisMinorClick(Sender: TObject);
begin
  Morris.Speed := Random(Morris.TopSpeed) + 1;
  Morris.Drive;
end;

procedure TForm1.btnDrivePorsche911Click(Sender: TObject);
begin
  Porsche.Speed := Random(Porsche.TopSpeed) + 1;
  Porsche.Drive;
end;

```

As you can see, test car objects are created (and set to an unparked state) when the form is created and destroyed when the form is destroyed. Each button sets the speed of the associated car to a random permissible value and then drives it. The results of this test are quite successful (see **figure 2**), but there are circumstances where the behaviour won't be as predictable and correct.

## Taking advantage of inheritance

One of the consequences of inheritance is that we get a class hierarchy and in that hierarchy are many branches. In any given branch there will be a base class and various descendants that are all supersets of that base class. Rather than treat instances of all these classes (of which there may be lots) using the correct types at all times, it can often be more convenient to treat them all like their common base class.

Our application involves a small class hierarchy branch with a common base class of **TCar**. If we wanted to represent lots of car objects we might declare an array of **TMorrisMinor** objects and another array of **TPorsche911** objects. Any operation we wanted to perform against all cars would need to be executed in two loops, one for each array. If any new **TCar** descendants were added to the hierarchy we would need to add a new array and go back to each place that we looped through the two arrays and add in another loop for the new array.

That would pose a maintenance problem, so it is usually better to declare a single array of **TCar** objects and just add all car objects to the single array. Anything inherited from **TCar**



Figure 2: A sample object executing a method.

can be treated like a **TCar** (as it is a superset of one) so it doesn't matter how many new classes are defined – the single array will suffice.

So we could add in a new data field (a **TCar** array) to the form, and initialise it when the form is created:

```
type
  TForm1 = class(TForm)
  ...
public
  Morris: TMorrisMinor;
  Porsche: TPorsche911;
  Cars: array[0..1] of TCar;
end;
...
procedure TForm1.FormCreate(Sender: TObject);
begin
  Randomize;
  Morris := TMorrisMinor.Create;
  Morris.Parked := False;
  Porsche := TPorsche911.Create;
  Porsche.Parked := False;
  Cars[0] := Morris;
  Cars[1] := Porsche;
end;
```

We can then add two more buttons to the form that share an **OnClick** event handler. If we set the **Tag** property of each button to identify an element of the array then the event handler can read the **Tag** of the pertinent button, then index into the array to access the **TCar** descendant.

```
procedure TForm1.btnDriveACarClick(Sender: TObject);
var
  Idx: Integer;
  Car: TCar;
begin
  Idx := (Sender as TButton).Tag;
  Car := Cars[Idx];
  Car.Speed := Random(Car.TopSpeed) + 1;
  Car.Drive;
end;
```

The idea is that either the Morris Minor or the Porsche will be represented by the **Car** object reference, so the appropriate top speed will be used as an upper limit for the random speed, and the car will then be driven.

## Polymorphism

Unfortunately the plan doesn't quite work out as you can see in **figure 3**. Despite pressing the new button that should drive the Porsche, the Porsche driving message is not displayed (and neither is the Morris Minor's when you press that button). This is because we have not used *polymorphism*.

Unless otherwise instructed, the compiler will work out what the target address of a call to any routine or method at compile time from the information it has to hand (the source code). In this case the compiler finds a **TCar** reference and so translates the call to **Car.Drive** into a jump to **TCar.Drive**. This is the natural behaviour of a compiler and is sometimes called *early binding* (routine calls are bound to target addresses before the program is executed).

In order to allow this generic representation of objects to be a useful technique we need the correct implementation of **Drive** to

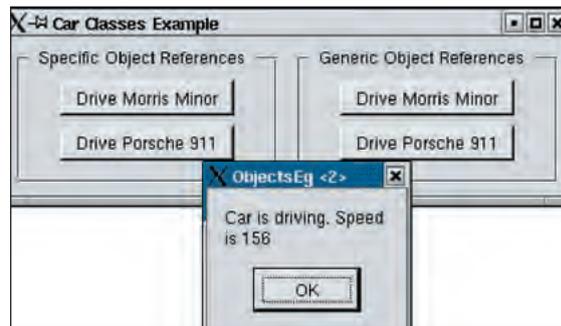


Figure 3: A lack of polymorphism impedes generic coding.

be executed, based upon which object is actually found at runtime. So if the **Car** object is actually a **TMorrisMinor** we need **TMorrisMinor.Drive** to be executed and the same applies for the Porsche. For this to work we need the compiler not to do early binding for this method; instead we need it do *late binding*, which means deciding which address to jump to at runtime.

Polymorphism (or late binding) allows you to call the same method on any object in a given branch of the class hierarchy (no matter how the object is being represented) and for the correct behaviour to ensue.

Since this behaviour is not the compiler's default, we need to specifically mark any methods that require polymorphism. When you add a new method to a class that needs to be polymorphic you must use the virtual directive. If you need to redefine a polymorphic method, to either change or extend it, use the **override** directive.

```
type
  TCar = class(TObject)
  ...
protected
  procedure SetTopSpeed(Value: Integer); virtual;
public
  constructor Create;
  procedure Drive; virtual;
  ...
end;

TPorsche911 = class(TCar)
public
  constructor Create;
  procedure Drive; override;
end;

TMorrisMinor = class(TCar)
public
  constructor Create;
  procedure Drive; override;
end;
```

So **SetTopSpeed** and **Drive** are marked as new polymorphic methods in **TCar**. The two descendant classes override the **Drive** implementation, but keep the inherited **SetTopSpeed** implementation. With these changes the program goes back to working just fine.

Note that the constructor is not marked polymorphic; polymorphism typically results from treating different *object instances* through the same type, such as **TCar**. A constructor is called through the class itself, not through an instance, so there is generally little point in marking it virtual. [LXF](#)

## About Brian Long

About Brian Long  
Brian Long is a UK-based freelance trainer and problem solver for Borland's *Kylix*, *Delphi* and *C++Builder* packages. His Web site 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 OOP support in *Kylix*. In the meantime, if there is something about *Kylix 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.



## WEB PROGRAMMING INTRO

# Practical PHP programming

## Paul Hudson starts a new series of tutorials on creating dynamic web content with PHP

**P**HP started as a quick Perl hack created by Rasmus Lerdorf back in late 1994, and since then, owing largely to its ease of use, flexibility, and sheer raw speed, the PHP module has climbed to the top of Security Space's online *Apache* module popularity report. This series is designed to take you from a rough beginner (some programming knowledge is assumed) to a competent user of PHP 4.2. Earlier versions of PHP may be incompatible with the code shown below, and we recommend you upgrade before trying the examples. Programmers coming from a C or Perl background should notice similarities between these languages and PHP. Note that PHP is case sensitive – **\$myvar** is different from **\$MyVar** – this often catches newcomers out!

### Getting Started

While PHP can be used straight from your terminal, the majority of users will want to use it through their web server. Nearly all Linux distributions come with some version of *Apache*, and it's certainly the most common web server out there, so I'm assuming you're using it also. If not, do not fret – PHP works almost anywhere you put it, and you can tailor any *Apache*-specific information to your own circumstances. See the box *Installing and configuring PHP for Apache* for more details on how to get PHP talking to *Apache*.

### Your first PHP script

Tradition dictates that your first program in a new language should output "Hello, World!" Many think that PHP stands for "People Hate Perl" (it's actually a recursive acronym for PHP: Hypertext Preprocessor), so we are going to break with tradition and start with this one-liner:

```
<?php echo 'Goodbye, Perl!'; ?>
```

Save that into the file "hello.php" in your public HTML directory and, once you have *Apache* and PHP configured correctly, view it in your web browser by entering the URL

```
http://localhost/hello.php
```

You should see the text "Goodbye, Perl!" in the browser. If not, read the box on debugging PHP programs to get help finding the problem. Let's break apart the one line to see what it contains. Firstly, we have **<?php**. This is the XML-compliant way to enter PHP mode; everything the PHP parser encounters between **<?php** and **?>** is read in "PHP mode" (as opposed to HTML mode), and is therefore treated as PHP code and executed.

Continuing on, **echo** is a pseudo function which prints data direct to the output. In our example, we print out "Goodbye, Perl!", but

we could equally have printed out variables, function return values, or other constants. Finally, we close our PHP code block with **?>**.

Everything PHP reads while not in PHP mode is sent straight to output – this is normally *Apache*, but could be STDOUT or even a file. There is no limit to the number of times you can switch to and from PHP mode. To illustrate this, here is another way to print out the same text, but this time we will switch into PHP mode and back out again.

```
<?php echo 'Goodbye; ?> Perl!
```

### Variables

In PHP you do not explicitly declare variables – they are created when they are first used. As a result, PHP variables are not fixed to one type – you can change an integer variable to a double, then to a string, and back to an integer again. We can see this implicit variable declaration in action in this following example. Two things to note: variables start with **\$** no matter what type they are, and **.** (a full stop) is the PHP string concatenation operator (it joins strings together).

```
<?php
$string1 = 'Goodbye, ';
$string2 = 'Perl!';
echo $string1 . $string2;
?>
```

### Getting external input

Previous versions of PHP automatically converted input from HTML forms into standard variables for you to use inside your scripts. However, recent versions of PHP have this feature disabled by default in order to strengthen the security of the standard install – input from users is now inherently distrusted. As a compromise between security and ease of programming, PHP creates several "super global" variables for you – these are variables from one variety of request input, eg HTTP **get** or **post**. PHP puts variables sent to you via HTTP get in **\$\_GET**, variables from HTTP post through **\$\_POST**, and a combined set of variables from **get** and **post** through **\$\_REQUEST**. Let's take a look at a simple form, called test.html, and see how PHP can be used to read the input.

```
<html>
<form method="post" action="form.php">
Enter your age: <input type="text" name="age" />
<input type="submit" />
</form>
</html>
```

Here is the companion PHP script also, which should be saved as form.php:

```
<?php
echo 'Your age is: ' . $_POST['age'];
?>
```

As you can see, the HTML page has its form action set to the PHP script, passing in just one field – “age”. This becomes `$_POST['age']` in our PHP script, which we echo back to the user. When you load your HTML page in your web browser (the URL will probably be `http://localhost/test.html` once you have copied the two files into your *Apache* public HTML directory), you will see the lone text box asking for your age. Try entering **25** as your age and submitting the form – you should see **Your age is 25** returned. Not very impressive yet, but keep reading!

## Using Arrays

Arrays are sets of data, stored in one variable. For example, you might have an array of seven values (one for each of the days of the week), or an array of football players in a team. Arrays in PHP look and act for the most part just like normal variables. You can create a new array in two ways:

```
<?php
$FirstArray = new array(); // explicit
// or
$SecondArray[] = "Blue";
$SecondArray[] = "Green"; // implicit
?>
```

Explicitly creating an array gives you an empty array – you need to add your items by hand. Notice that the variable itself, `$FirstArray`, is still implicitly declared – we just explicitly define it as an empty array. Implicitly creating an array, as we see in the second option, is where we create an array by adding its first item with the `[]` operator. You can dynamically add as many variables as you like (strings, integers, other arrays, etc) by using the `[]` operator. In our example, `$SecondArray` contains two elements of data – “Blue” at index 0, and “Green” at index 1. These can be accessed with `echo()` by using `$SecondArray[0]` and `$SecondArray[1]` respectively. You can make your life easier by using associative arrays – that is, rather than using 0 and 1 as your indices, you could use `foo` and `bar`, or `JacksFavouriteColour` and `JillsFavouriteColour`. Compare `$MyArray[13]` to `$MyArray['username']` – I think the latter is much easier to read and remember!



## Adding more functionality

Two points commonly cited by users as to why they program in PHP rather than other languages are: it has incredible database capabilities (over fifteen different database systems supported at the time of writing!), and it has a large collection of non-trivial functions available to users. We are going to touch on databases briefly later on, but for now let's take a look at some of the basic functions PHP offers to make your scripts more powerful. To start with, `isset()` is a commonly-used function which returns true if a variable is set in the current context. `set` does not necessarily mean the variable has a value – a value may well be blank if a form field was left empty, for example. We can rewrite our age script into one page, `form.php`, by using `isset()` like this:

```
<html>
<?php
if (isset($_POST['age'])) {
    echo "Your age is: " . $_POST['age'];
} else {
?>
<form method="post" action="form.php">
Enter your age: <input type="text" name="age" />
<input type="submit" />
</form>
<?php
}
?>
</html>
```

Putting HTML forms and the PHP code to handle them into one page is very common practice – above you can see that we are checking our `$_POST` super global to see whether age is set, and printing out our HTML form if it is not. Try running the page – you will notice it operates just the same. `time()` is a simple function which returns the number of seconds since the Unix epoch (January 1st, 1970, 00:00:00 GMT). When combined with `date()`, you can convert your timestamp to just about

Your first form should look like this.



## Installing and configuring PHP for Apache

Get started with the minimum of fuss

Download PHP from <http://uk.php.net/downloads.php> and extract it. Provided you have *Apache* installed with Shared Module support, you need only do the following:

1. `cd /path/to/php-xxx`
2. `./configure --with-mysql --with-apxs=/path/to/apache/bin/apxs`
3. `make`
4. `make install`
5. `cp php.ini-dist /usr/local/lib/php.ini`
6. Edit your `httpd.conf` or `srm.conf` file and

check that these lines are present and not commented out:

```
AddType application/x-httpd-php .php
LoadModule php4_module
libexec/libphp4.so
```

7. Restart *Apache*. This command varies system to system, but `/path/to/apachectl start` should work.

If you do not have *Apache* installed already, please refer back to our detailed instructions in *LXF 26*.

For a complete list of PHP's configure options, read <http://uk.php.net/manual/en/install.configure.php>. I do not want to give too much away about future articles in this

series, but I would recommend you at least use a few of the following configure options:

```
--enable-inline-optimization
--with-gd=/path/to/gd
--with-jpeg-dir=/path/to/libjpeg
--with-png-dir=/path/to/libpng
--with-t1lib[=DIR] --with-pdfcpl[=DIR]
--with-ming[=DIR] --with-xml
--enable-xslt --with-xslt-sablot
```

If you are running on a 64-bit processor, remove `--enable-inline-optimization` as *GCC* PHP compilation has been known to fail in 64-bit environments. Remember: it is always best to read the *INSTALL* files bundled with the PHP source for more detailed instruction on installation and configuration.

any format imaginable. **date()** takes two parameters – a string made up of special character codes representing the format you want to output, and a timestamp. Here are some examples to get you started:

```
<?php
$ShortDate = date("j/n/y", time()); // returns eg 25/12/2002
$LongDate = date("F jS, Y", time()); // returns
eg December 25th, 2002
$LongDateTime = date("g:ia F jS, Y", time()); // returns
eg 11:39am
December 25th, 2002
?>
```

Two functions which are very useful are **explode()** and **implode()**. **explode()** takes a string with a separator (e.g., split by commas), and creates an array of the individual elements. That is, "this,is,a,test" would become an array with four elements – "this", "is", "a", and "test". **implode()** does quite the opposite – our four-element array can be imploded into a string with separators of your choice. So, to turn the string "comma,separated,string" into "commalseparated!string" we could use code like this:

```
<?php
$string = "comma,separated,string";
$exploded = explode(",", $string);
$imploded = implode("!", $exploded);
?>
```

**phpinfo()** is a very popular function owing to it being so useful when debugging. Given no parameters, **phpinfo()** gives you all the information about your current PHP configuration – where php.ini (the configuration file) is, what web server it is running on, what PHP modules you have compiled in, the configuration settings that are in place, as well as all the predefined variables available to you. I cannot stress enough that **phpinfo()** is incredibly valuable as a debugging tool, as it gives you the cold, hard facts about your configuration and lets you rule out obvious problems immediately.

Another great debugging function is **var\_dump()**. **var\_dump()** takes one or more variables as parameters, and returns neatly-structured information about those variables. Arrays are traversed fully, with all their variables defined and printed out. I suggest you encase your **var\_dump()**s in HTML **<pre>** elements so that the output is shown properly formatted.

A helpful function for use with arrays is **count()** – it takes an

array as its sole parameter, and returns the number of elements in that array. Remember that PHP uses zero-based arrays, so the last element in your array can be referred to as

**count(\$myarray) - 1**.

Finally, before we move onto a complete example of PHP at work, I want to introduce you to the **mail()** function quickly. **mail()** takes a minimum of three parameters – the email address to send to, the subject line to use, and the message to send. You can also specify a fourth parameter if you want to add your own extra headers. Text in your message, and also individual headers, should be separated by a new line (**\n**).

```
<?php
$message = "Line 1\n Line 2\n Line 3\n Line 4\n Line 5\n";
mail("john@dunbar.com", "Testing PHP mail()", $message,
"From:
jack@jill.com");
?>
```

By specifying special headers in the fourth parameter, you can send HTML mails, specify the Reply-To address, or add any other special information you choose.

## A complete program

You should now know enough PHP now to put together a fairly good PHP script. No databases just yet, but never fear – we will get there soon enough!

```
<?php
if (isset($_POST['user'])) {
if (($_POST['user'] != "") && ($_POST['pass'] != "")) {
echo "We have successfully created a new account for
you!";
mail($_POST['email'], "Account Information", "Thanks for
creating an
account at somesite.com!", "From: accounts@somesite.com");
} else {
echo "You must enter a username and password!";
}
} else {
?>
```

To create an account, enter a username, password, and email address below.

```
<form method="post">
Username: <input type="text" name="user" /><br />
Password: <input type="password" name="pass" /><br />
Email Address: <input type="text" name="email" /><br />
<input type="submit" />
</form>
```

```
<?php
}
?>
```

## Database interaction

Ever since version four of PHP was released, the PHP community has benefited from having a base library of *MySQL* installed by default – built directly into the source code. This was the direct result of many meetings between the PHP and *MySQL* developers to try to figure out how to best aid the already-strong relationship between the two programs.

As a result of such helpful integration, and, of course, because *MySQL* is an excellent database manager on its own merits, *MySQL* is by far the most popular database system for use inside

## Getting help with PHP

There are many options to choose from

PHP has all the usual fast and friendly community support we have all come to know and love – very active mailing lists, a popular IRC channel on EfNet where the maintainers hang out (#php), and a solid collection of online tutorials.

Add to these its first class online manual which is packed with user comments and you really cannot go wrong. Check <http://uk.php.net/support.php> for a full list of PHP support options.

If you are interested in

buying a book, I recommend you consider either *Web Application Development with PHP 4.0* (New Riders, 2000) by Ratschiller and Gerken, or the *PHP Developer's Cookbook* (Sams, 2002) by Sterling Hughes, as both books offer very practical advice that can often be put to use immediately.

Of course, you are always welcome to write to the *LXF Answers* team, and we will do our best to solve your problem for you! :)

## Debugging PHP programs

### Straightforward fault-finding tips

Debugging, as most of us know, is an art that is best learnt over time with experience. To give you a headstart, however, here are answers to some frequently-encountered problems.

**1** I see the sourcecode to my script output

This normally occurs when you have misconfigured *Apache*. Check that you have followed my installation guidelines properly.

**2** I get a "parse error" message

Check that you have entered the script exactly as you see it – PHP is very specific about its syntax. Common mistakes include forgetting to end each statement with a semi-colon, missing off closing quotes from strings, forgetting to start variables with a dollar sign, and the like – try reading your script thoroughly to catch these little errors.

**3** My browser cannot find the server "localhost"

This is an *Apache* error – ensure that you have started your server on port 80.

**4** I get a 500 server error from *Apache*

If you see this as a result of trying to load a PHP script, then it is quite possible that the page you are loading does not exist. Failing that, it is almost certain that you have grossly misconfigured either *Apache* or PHP!

**5** I cannot track my problem down

Remember that *Apache* frequently stores information about problems in its error\_log file – very often you can find your elusive problem answered here!

**6** I see `$String1 $String2` rather than **Goodbye, Perl!**

This is a common problem when you are getting started, and is down to how PHP handles quotes. All strings surrounded by single-quotes (') don't have their variables replaced by their values – that is, when you type `$foo = '$bar'`, PHP assumes you actually want a dollar sign and the word "bar" assigned to your variable. `$foo = "$bar"`, which uses double quotes ("), forces PHP to replace `$bar` with the value of the variable `$bar`.

PHP. In order that you can follow this code, you will need to download *MySQL* from [www.mysql.com](http://www.mysql.com), and follow their installation instructions. Create a user called "php", and give it a password – note that I have used the password "alm65z", but you will probably want something different. I will be giving a thorough *MySQL* installation guide in a subsequent tutorial, but until then you will have to rely on the official documentation.

```
<?php
mysql_connect("localhost", "php", "alm65z");
mysql_select_db("phpdb");

$result = mysql_query("select * from users where id = '$ID'
and password = '$pass;");
if (mysql_num_rows($result)) {
echo "You have been authenticated successfully!";
} else {
echo "Invalid username and password!";
}
?>
```

Our first two lines connect to the *MySQL* server on localhost with username "php" and password "alm65z", selecting the database "phpdb" for querying. Running a query on *MySQL* is as simple as using the `mysql_query()` function. It returns a reference to all the rows which matched our query as opposed to the data itself – we get that later on. Once we have our reference, we can run various functions on it to get more information about the data – in the example, we call `mysql_num_rows()`. This function returns the number of rows returned by the query and, as it is inside an `if` statement, the `if` statement evaluates to true if we have at least one row returned. PHP supports a wide variety of databases, including *PostgreSQL*, *Oracle*, and even Microsoft's *SQL Server*. You can run a query on *PostgreSQL* simply by changing `mysql_query()` to `pg_query()`, and similarly you can check the number of rows returned by `pg_query()` with the `pg_num_rows()` function. Let's take a look at one final database example for this month – a password reminder.

```
<?php
if (isset($_POST['email'])) {
mysql_connect("localhost", "php", "alm65z");
mysql_select_db("phpdb");
```

```
$EmailToCheck = $_POST['email'];
$result = mysql_query("select username, password from
users where emailaddress = '$EmailToCheck;");
if (mysql_num_rows($result)) {
    $r = mysql_fetch_array($result);
    $TheirUser = $r['username'];
    $TheirPW = $r['password'];
    mail($EmailToCheck, "Password Reminder", "Your
username is $TheirUser,
and your password is $TheirPW", "From:
password@somesite.com");
    echo "Your info have been emailed to you at
$EmailToCheck";
} else {
    echo "We do not have an account for that email address!";
}
} else {
?>
```

To recover your username and password, enter the email address you used to register with us.

```
<form method="post">
Email Address: <input type="text" name="email" /><br />
<input type="submit" />
</form>
<?php
}
?>
```

The new parts in the example above are around the `mysql_fetch_array()` function. Owing to the fact that one *MySQL* query may return many rows, PHP lets you handle each row individually – `mysql_fetch_row` takes the next row (beginning with the first one), reads the values of all its fields, and places them into an associative array (`$r`, in our example). So, it automatically sets `$r['username']` and `$r['password']` to be the values of `username` and `password` for the row which matches the email address supplied.

## Conclusion

This tutorial should have given you a kick start in the world of PHP – you have got enough information now to continue on and create a variety of simple scripts to liven up your website. [LXF](#)

## About Paul Hudson

Paul Hudson is a London-based web developer specialising in PHP and Perl. He can be emailed at [jedihudzilla@hotmail.com](mailto:jedihudzilla@hotmail.com).

## NEXT MONTH

Next month we will look at a variety of helpful PHP routines to handle authentication over HTTP, uploading files through HTML forms, and more. If there is something you want to see covered here, drop us an email and we will try our best to incorporate it into a future instalment.

# 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*...



## Configure question

**Q** I began buying *Linux Format* mags almost from the first issue but have only recently set up a PC running SuSE Linux 7.0 – so I'm a beginner.

I decided to load a game from the May 2002 DVD (e.g. *Attal*) just to see whether I could manage it.

Following the general instructions on p105, I managed to get the files unpacked in a folder. *cd'd* into that folder and examined README ("see INSTALL" – very useful!! But I suppose necessary for beginners like me)

Typed `make -f Makefile.dist` at the *bash* prompt and got the error message:

```
aclocal
make: aclocal: Command not found
make: *** [all] Error 127
```

Can only guess that *make* was looking for a file called *aclocal* and didn't find it – I'm not sure where it is either!

This is only one example of a problem I've had installing a game, as I tried to install at least half a dozen other games afterwards with no success either.

Oh yes, and by the way, I like *Linux Format* – especially the *Kylix* tutorial series which I went through on a PC running '98 and *Delphi*.

Will try to load *Kylix2* soon – if I can.

Kind regards

**Hugh Halpin**

**A** *aclocal* is distributed as part of the *autoconf* package, which is used to generate the *configure* script which many developers use to assist people who are building their source code. If you're missing *autoconf*, then it's quite likely that you're missing other C development packages which you need to build source code. SuSE may have an option in *YaST* to install all C

development packages, but you generally need *autoconf*, *automake*, *libtool*, *glibc-devel*, *gcc*, *binutils* and any other development package for libraries which the program will be linking to. For example, if it uses *zlib*, then you will need to install *zlib-devel* which provides the C headers for the library so that programs can compile against it.

## Lost scanner

**Q** The other day I upgraded my Red Hat 7.2 workstation. I used to work with *XSane-0.82-3* and *xscanimage* without problem. Now I have upgraded to Red Hat 7.3 (created my own CDs) and *XSane* appears to be upgraded as well to version 0-84. The commands `sane-find-scanner` and `modprobe g_NCR5380`

`ncr addr=0x280 dtc 3181e=1` work well, as they did before. But when I start *XSane* or *xscanimage* both report no scanner has been detected. How is this possible? The `soft-link /dev/scanner` to `/dev/sg0` or `/dev/sga` is still there.

Thanks

**Huub**

**A** As you can load the the module, it would suggest that the kernel can find a device at the appropriate I/O address. After you load the module, if you run `dmesg` from the command line, it will list all of the kernel messages which it has sent to the *syslog* daemon. You should notice a line or two for the loading of the module, such as the version and copyright information, but it should also produce some information identifying the device it has found and, hopefully, can tell you if it thinks it is functional or not. It should also tell you which device, or major/minor, the device belongs to, so you can make the appropriate symlink. Does

```
sane-find-scanner -v
```

tell you which device the scanner is on, and identify it correctly?

You may want to look at downgrading your copy of *XSane* to 0.82, to see if you experience the same problems. At the time of writing, the latest release of *XSane* is currently 0.86, so if there are RPMs of that available, it may be a desirable upgrade.

## ToUCam troubles

**Q** I am trying to get my Philips ToUCam (Model 730) to work under SuSE 7.3 (2.4.10 kernel and *XFree 4.2.0*). Getting the kernel to recognise the camera is not a problem and *usbview* correctly reports its presence. However, whenever I start up *xawtv* I get the message "no overlay support" and my system locks up with the caps-lock and scroll lights blinking (a kernel panic I assume).

I tried a freshly downloaded version of *Gnomemeeting* and as soon as I checked the preview option the same thing happened. I have even tried telling *xawtv* not to use overlay and still it locks up my computer.

I keep looking back at the excellent USB article by Simon Goodwin (LXF22) and marvel at how he got his webcam to work. As far as I know I have done everything I should, including the



The *pwc* driver for Philips webcams can be accessed using many different apps, including *xawtv*.

purchase of a Linux compatible webcam and I still get no joy.

My hardware setup is a 1.4GHz Athlon with 512MB RAM and a Matrox Millennium G450 (16MB DDR) attached to an Iiyama VM Pro 413.

When camera is loaded I get:

```
root_22> lsmod
Module Size Used by
pwc 34880 0 (unused)
videodev 4608 1 [pwc]
evdev 4160 0 (unused)
input 3264 0 [evdev]
usb-uhci 20912 0 (unused)
usbcore 52256 1 [pwc usb-uhci]
```

I would greatly appreciate some help in getting my webcam up and running under Linux. Thanks

Martin Knott  
Edinburgh

Firstly, we would guess that the kernel panic is independent of the overlay

issue. Can you check your system logs, or use the **SysRq** commands to produce any useful debugging information? Certainly, your kernel needs updating, as there have been marked improvements to the *pwc.o* driver since 2.4.10 was released. You may also want to install the *pwcx.o* module, which provides compression support for the camera. The latest release of the *pwc* and *pwcx* modules are available at [www.smc.demon.nl/webcam](http://www.smc.demon.nl/webcam).

While the camera is supported by Linux, it may be that there was a specific bug in the 2.4.10 kernel, either in the USB code or specifically in the *pwc.o* module, which causes the kernel to fail when you use the camera.

The overlay support error is due to the camera not supporting the *Xvideo* extension. This is used by *video4linux* when writing data back to a video

device, and since the camera does not accept any input, then it is not going to support any overlays. Try adjusting the *xawtv* settings file to default to capture mode.

## PCTel module

Hi, I recently decided to go 'Linux' as I was fed up with Microsoft's constant 'buy Microsoft' campaign, and chose Mandrake Linux 8.2. I have tried to get my modem working but to little avail. Yes, before you wonder, it is the PCTel Winmodem and I'm having an agonising time getting it to work so I would very much like your help.

My Kernel version is 2.4.18 and, like I have previously stated, I own a PCTel Winmodem. I have managed to do the following:

- 1 Extract the tar file,
- 2 Run the *make* command

which, in reply, gave me **Nothing to be done for all targets**

3 Change to su

4 **mkmod /dev/ptel c 62 79**

5 **ln -s /dev/ptel /dev/modem**

which, in reply, gave me

'**dev/modem**': File Exists

6 **ln -s /dev/ptel /dev/ttyS15**

7 **chgrp uuap /dev/ptel**

8 **chmod 666 /dev/ptel**

9 **cp ptserial.o**

**/lib/modules/2.4.18mdk/misc**

which, in reply, gave me **cannot create regular file '/lib/modules/2.4.18mdk/misc/ptserial.o': no such file or directory**

I then proceed to **insmod -f ptserial.o** anyway and it says the following:

```
Kernel Mismatch ptserial.o was
compiled to run on kernel 2.4.0
while you are currently running
2.4.18-6mdk.
```

Do I have to recompile the .o file – if so how?

If you think of other things to do please show me exactly how as, like I say, I have just moved to Linux recently.

Thanks Very Much,  
Dan

In later releases of the 2.4.0 kernel, the directory structure of `/lib/modules/<version>/` has changed considerably, and you probably want to put your module in `/lib/modules/2.4.18mdk/kernel/drivers/net/`. You will either need to recompile the module, or obtain a different version of it, in order to use it with the 2.4.18 kernel from Mandrake. There are a number of different drivers available for PCTel modems, and you should logically compile and install them in turn in order to find out which one works correctly with your system.

Rather than repeat it all here, you should take a look at <http://pctelcompdb.sourceforge.net/ptel.html>, which has documentation and the locations of the downloads for the various drivers. It is slightly more complex than simply loading a module, as you need to use *lspci* to find out what chipset the modem has, although after that the modem should work.

## CGI printenv

The Linux in my system is Mandrake 8.0 and I was following your tutorial on CGI programming article in which

## A QUICK REFERENCE TO: kernel compilation

Generally once Linux has been installed, many people like to head over to [ftp.kernel.org](http://ftp.kernel.org), and download the latest kernel release. This ensures that they are running a system with the latest bug fixes, which may affect stability and security, as well as providing any additional device support for specific hardware which their system has. Many distributions come with kernel sources, which make it easy to add or remove support for certain things, but if you are wanting to apply third-party patches to the kernel, or otherwise change it, you will need to ensure that you have a version of the kernel which is clean, otherwise the patches will fail.

Rebuilding the kernel has a number of stages, which involve setting up the kernel with the appropriate capabilities, compiling the code for the kernel, and then installing both the kernel, and any modules, on the system. The following stages will work equally well with either a kernel downloaded from [ftp.kernel.org](http://ftp.kernel.org), or the source code for a distro's kernel which can generally be found on their CD or ftp site.



When building a new kernel, it is important to include all the capabilities we need to make our system run properly.

The kernel configuration is stored in `linux/.config`, which can be copied out and saved for use in a later kernel. If we store our kernel's configuration elsewhere, we can copy it back into our kernel tree, as `linux/.config`, and then run;

```
make oldconfig
```

which will step through the existing `.config`, and will prompt us for any new configuration options which will have been included since we last built our kernel. This equally works if we patch our existing kernel, and want to know what new things it can do.

If we want to reconfigure our kernel, we can use **make config**, **make menuconfig**, or **make**

**xconfig**, which use different user interfaces for configuration.

**menuconfig** should be used for a console system, or **xconfig** for a box with an X server running.

Once our kernel is configured, we can build and install in one go;

```
make dep && make clean &&
make bzlilo && make modules &&
make modules_install
```

This will clean up the kernel tree, removing old binaries, build the new one, then install it with *Lilo* and install the modules. **make bzlilo** installs the kernel in `/vmlinuz`, and runs `/sbin/lilo` – if your kernel is stored in `/boot`, you will need to copy it over manually and rerun `/sbin/lilo` manually.

## FREQUENTLY ASKED QUESTIONS: XFree86

### FAQ I'm trying to configure XFree86 4.0, but there is no XF86Setup program. How do I setup XFree86?

XFree86 uses a single binary, XFree86, for everything, including starting the X server, and building the configuration script. You can configure X with;

```
# XFree86 -configure
```

which will autodetect your video card, and prompt you for other information, including monitor specifications and card capabilities which it is unable to detect. Once it has built the configuration file, you should be able to run **startx**, which will start up the X server, then you can test the configuration and adjust it as appropriate.

### FAQ My screen is appearing at 640x480, and the desktop scrolls around. I know my monitor is capable of 1024x768, but X always uses 640x480.

Usually when X is setup, it will include all of the resolutions which the system will run at within the configuration. If you can run at 1024x768, then X will be setup for 640x480, 800x600 and 1024x768. However, it will first attempt to use 640x480, but with a virtual resolution of 1024x768. You can change resolution using **Ctrl-Alt-Numpad+** and **Ctrl-Alt-Numpad-**, but if you want to use 1024x768 as default, you will need to edit the **Screen** section of `/etc/X11/XF86Config`, and have it list 1024x768 before any other resolutions.

### FAQ I'm trying to run programs from my laptop on my workstation, but X applications won't work with telnet. How can I make them appear on my workstation?

There are two ways to run X applications over a network. The simplest is to install **SSH**, which does X11 port forwarding, so you can execute programs on the remote shell and have them connect to the local X server. Alternatively, you can use **xhost** to permit the remote machine to

connect to your X server, then point your client at it. To allow the machine 'bob' to connect to our local machine 'fred', we do;

```
$ xhost +bob
```

on 'fred', then telnet, rsh or ssh to 'bob' and do;

```
$ export DISPLAY=fred:0.0
```

Once **\$DISPLAY** is set, you can run any X applications, and they will appear on the X server running on 'fred'.

### FAQ Does X support multiple monitors and more than one video card?

XFree86 supports multiple monitors using the **Xinerama** extension. This allows the X desktop to span across monitors which may be connected to one, or more, video card. Each video card requires its own **Device** definition with XF86Config, supplying its PCI or AGP information, so it can identify individual cards. Setting up **Xinerama** requires some knowledge of the hardware configuration of your machine, so you may want to look at the **Xinerama HOWTO** at [www.linuxdoc.org](http://www.linuxdoc.org).

### FAQ My DVD player complains about the Xv extension being missing from my X server. What is this, and how do I fix it?

The **Xvideo** extension is used to improve throughput of motion video, by pushing the data straight at the video card, rather than having X process it, which reduces CPU load. Not all video cards support Xv, as it requires specific hardware, so if your machine does not have it, then there isn't too much you can do. Most programs which use Xv also support regular X11 output, so you should be able to reconfigure the application. However, it will be considerably slower than using **Xvideo**.

### FAQ My mouse has a wheel on it, so how do I make programs use it?

Within the **InputDevice** section for the mouse, you need to include;

```
Option "Buttons" "5"
```

```
Option "ZAxisMapping" "4 5"
```

Which will then map the wheel

« you had a HOWTO configuration of Apache. The Apache server is running OK on my system as I get the web server default page. But I am unable to obtain the list of environment variables when I look for `http://localhost/cgi-bin/printenv`.

I have tried to locate the **cgi-bin** file as suggested but was unable to find it on Mandrake. I have found the `http.conf` but in a different location from that suggested in the article (I suppose each distro has their own format for filing system). However, I am not able to find **cgi-bin** and therefore not able to configure Apache as recommended.

I have been a casual Linux follower since 1996, but to date I find your magazine being the best informative publication. Keep up the good work.  
Anvar Alizadeh

**A** You may not have a **cgi-bin** directory on your system, as different distributions install Apache in different places, or with different configurations. Usually **cgi-bin** lives in `/usr/local/apache/`, but some systems have `/home/httpd`, and other directories. Once you find your **htdocs** directory, make a **cgi-bin** directory in the same place, then edit the `httpd.conf` file, where it has the **DocumentRoot** entry, to use a **ScriptAlias** for the **cgi-bin** directory which you just created;

```
ScriptAlias /home/httpd/cgi-bin /cgi-bin
```

You also need to create a **<Directory>** section in `httpd.conf` for **cgi-bin**, to allow CGIs to run within the directory;

```
<Directory "/home/httpd/cgi-bin">
```

```
AllowOverride None
```

```
Options None ExecCGI
```

```
Order allow,deny
```

```
Allow from all
```

```
</Directory>
```

Obviously, in both cases, you will need to edit the path to **cgi-bin** depending on the way Apache has

been installed. If you think you might already have a **cgi-bin**, then you can look for the **printenv** script with

```
locate printenv
```

which should not be confused with

```
david@macha:~ (pts/8)
david:~/hine-1.0-8.0_104/./configure: more
creating cache ./config.cache
checking host system type... i686-pc-linux-gnu
checking target system type... i686-pc-linux-gnu
checking build system type... i686-pc-linux-gnu
checking for a BSD compatible install... /usr/bin/install -c
checking whether build environment is sane... yes
checking whether make sets $(MAKE)... yes
checking for working aclocal... found
checking for working autoconf... found
checking for working automake... found
checking for working makeinfo... found
checking for strerror in -lcposix... no
checking for gcc... gcc
checking whether the C compiler (gcc) works... yes
checking whether the C compiler (gcc) is a cross-compiler... no
checking whether we are using GNU C... yes
checking whether gcc accepts -g... yes
checking how to run the C preprocessor... gcc -E
checking for ANSI C header files... yes
checking whether make sets $(MAKE)... (cached) yes
checking for a BSD compatible install... /usr/bin/install -c
--more--
```

The **configure** script tests for different packages, including **gcc**, and will fail if they are missing.

movements to buttons 4 and 5 on the mouse. You can then use a user-space utility, such as *imwheel*, to have these events appear to client programs as **PageUp** and **PageDown**, or indeed any combination of key presses which the client will make use of.

**FAQ** I have a crazy keyboard with a load of extra buttons. Can I use them to control things like *xmms* and *xscreensaver*?

Using *xev* to capture the keycodes, we can create a *.xmodmap* file, which is used by *xmodmap* to assign key names to our keys. For example, on a Microsoft Natural Keyboard, with the Internet keys, the **Sleep** button is number **223**, which we can map to a **F** key;

**keycode 223 = F33**

We can then use our window manager's key binding application, to have it execute a specific command, such as **xscreensaver-command -lock**. The use of special keys was covered in *Linux Format* 28.

the *printenv* binary which is part of the GNU *sh* utilities package. The *cgi* version of *printenv* is a Perl script, and is clearly identified with `#!/usr/bin/perl` as the first line.

## Dependencies

**Q** I have just recently bought and successfully installed SuSE Linux 8.0 and love it. My problem lies here, with RPM's, *tar.gz* and Dependencies.

Every time I download an RPM file, or *tar.gz* for that matter, it never compiles – on RPM's it goes right down to the most smallest file like *libglcore.so.1* which I cannot find anywhere for a start.

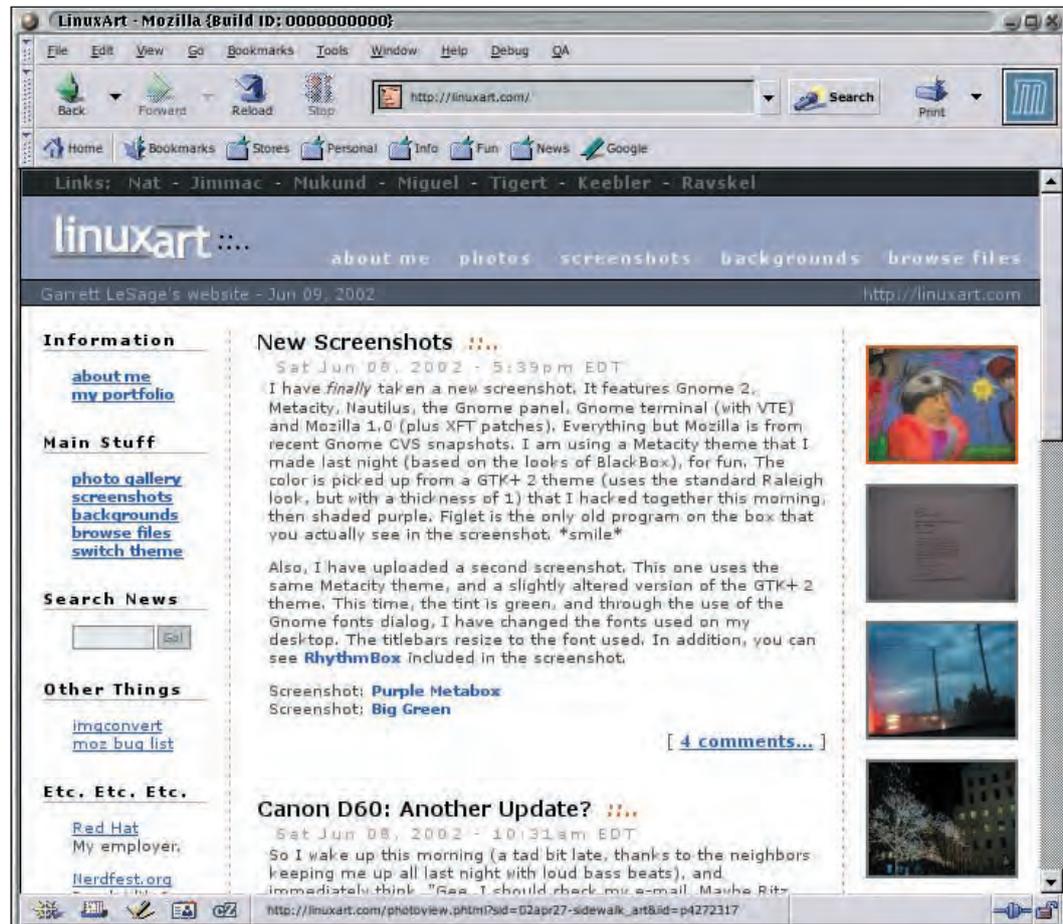
When I try *tar.gz* it says the following at the *bash* prompt

Checking for Mawk... No

Checking for Gawk... No

Checking for GCC... No

Checking for CC... No



LinuxArt.com uses many features of Mozilla, including external style sheets.

Checking for CC... No

Checking for CL... No

Checking whether make sets

`$(MAKE)...` No

Configure: Error: No acceptable CC found in \$Path

I am running an AMD Athlon 1GHz, 256MB of RAM and 120GB of hard disk space.

Where can I find these files, how can I install them and how can I ensure myself easier package installation in the future?

Many Thanks  
Craig

**A** The error you are experiencing is because you are missing a C compiler on the system. The *GNU C Compiler*, or *GCC*, is the standard compiler on Linux, and will be distributed on the CDs which came with your SuSE 8.0 distribution. You may find, however, that there are other things missing from your system, and you will need to install them one by one as the *configure* script throws up the appropriate error message.

The RPM errors are because the RPM you are trying to install depends

on a specific file, which is contained within another RPM. In the case of *libglcore.so.1*, which is distributed as part of the *Mesa* package, you will need to go and install *Mesa* before attempting to install whichever RPM you are trying to use. If you get stuck, and are not sure which package provides which specific files, there is a search tool on [www.rpmfind.net](http://www.rpmfind.net), which will allow you to insert a dependency, such as *libglcore.so.1*, and will then proceed to tell you which package provides that. However, you should be aware that other distributions use RPMs too, such as Red Hat and Mandrake, and you should generally only install SuSE packages, as the locations of certain files may be different from what RPM expects if the RPM was built on a RH box.

## Missing configure

**Q** I am running the Mandrake 8.2 distro from the May 2002 issue of your excellent mag. I have installed several .rpm binaries successfully but can never get tar files to work. I followed the instructions carefully

and managed to unpack the .tgz file for *K-3D* to my home directory. I cd to the directory and type *.configure* as instructed, but all I get is:

**bash: ./configure: No such file or directory**

I have tried this in my user login and as root and get the same message in both cases. I also seem to be missing the *make* and *make install* commands as well.

Am I doing something wrong or is my Mandrake installation at fault?  
Kev Learwood

**A** The *configure* script builds the Makefile used by *make*, so if *.configure* is missing, or fails, then the *make* commands will not work. If the *configure* script is not there, then you need to use *autoconf* to build it. As *K-3D* is still heavily in development, you may wish to check out <http://k3d.sourceforge.net/>, and look for an up-to-date tarball or RPM of the program.

Being missing a *configure* script is fairly uncommon, so if you are still having issues with source distributions, then you may wish to post a detailed

error log on the *LXF* forum, so people may be able to assist, and help you figure out what extra packages need to be installed.

## Lacking style

**Q** A quick question on the version of *Mozilla* you distributed on last month's coverdisc.

I don't know whether I'm doing something wrong, but I can't get it work with external Style Sheets. It will use embedded and inline styles OK but won't import an external one.

The syntax I'm using is

```
<link rel="stylesheet"
type="text/css">
```

as you would expect.

Confusingly, the Stylesheet is correctly picked up in my currently installed *Netscape 4.75*. *Mozilla* doesn't recognise it at all.

If it makes any difference, I'm running *Red Hat 7.0* (and *fluxbox*, thanks to your review!) Can you offer any advice?

Great magazine by the way – it's helping me to develop my Linux knowledge bit by bit.

**Peter Harrison**

**A** *Mozilla* does correctly support CSS, versions one and two, so it should load a stylesheet from an external source. The appropriate syntax, according to the *Mozilla* CSS documentation is;

```
<link rel="stylesheet"
href="/style.css" type="text/css">
```

If you're still experiencing problems, then it's quite probably a bug in the pre-release of *Mozilla* which you are running. As *Mozilla 1.0* has been released, you may wish to download the latest build from the Internet (or check our CDs), and check to see if CSS works correctly with a more up-to-date version.

## SCSI CD

**Q** I have just got a new computer and decided to try not to use MS products, so I had *Red Hat 7.2* loaded for me and purchased the O'Reilly *Learning Red Hat Linux* book, but I now live in Bangkok where I have your magazine sent out to me when I'm not travelling to the UK so cannot buy it for myself.

I purchased my first program *Corel WordPerfect 8* for Linux and have been unable to load it. I think

```

david@macha:~ (pts/6)
Reformatting resize_reiserfs.8, please wait...
RESIZE_REISERFS(8)          RESIZE_REISERFS(8)

NAME
  resize_reiserfs - Reiserfs filesystem resizer

SYNOPSIS
  resize_reiserfs [-s [+|-]size[CKMG]] [-j dev] [-fqv]
  ] device

DESCRIPTION
  The resize_reiserfs tool resizes an unmounted reiserfs
  file system. It enlarges or shrinks an reiserfs file sysi
  tem located on a device so that it will have size bytes or
  size=old_size +(-) size bytes if the + or - prefix is
  used. If the -s option is not specified, the filesystem
  will be resized to fill the given device. The size param
  eter may have one of the optional modifiers K, M, G, which
  means the size parameter is given in kilo-, mega-, giga-
  bytes respectively.

  The resize_reiserfs program does not manipulate the size
  of the device. If you wish to enlarge a filesystem, you
  
```

**ReiserFS allows you to resize a mounted file system, as long as there is space on the partition.**

I may know the problem and need your help to confirm and ask you for advice.

I had a look under "user mount tool" to see if the CDROM had been mounted or not and found the following

Directory	Device	Type
/mnt/cdrom/	dev/cdrom	iso9660

I then clicked the Mount check box and got an error message reading **Mount:No medium found. Why?**

So I had a look at the hardware browser and found the following:

CD-ROM Drives:-	Lite-on LTR	24102B,	under device information:
"manufacturer"	Unknown	"Driver"	ignore
Device	/dev/scd0		

Have you got any suggestions? I really don't want to fall at the first hurdle, and I'm the only person using Linux in Thailand and really want to master and promote it once I'm up and running, so any help would be very much appreciated

**Tony Ryder**

**A** /dev/scd0 is used to identify SCSI CDROM drives, either those connected to a physical SCSI host adaptor, or one using the *ide-scsi* module, which is commonly used with CD burners or other devices requiring SCSI emulation. In order to mount /dev/cdrom, you will need to be sure that /dev/scd0 is symlinked to /dev/cdrom, which you can do with;

```
# rm /dev/cdrom
# ln -sf /dev/scd0 /dev/cdrom
```

If you are still experiencing problems, use *dmesg* to check the kernel message log, and see which device it is trying to read from. Of

this as I hope some time in the future to rid myself of Windows 2k partitions and use *Codeweavers Crossover Office*, that you recently reviewed instead ;-)

**Brian Durant**

**A** We think you have the wrong end of the stick in order to increase the size of your partition and filesystem. Using up the extra space is a two stage process, involving both Linux and Windows for *Partition Magic*. You should reboot into Windows, then use *PM* to increase the size of the Linux partition to take up the whole of the disk. If it complains about the *ReiserFS* filesystem, have it resize the partition without touching the filesystem on it. Once this has completed, you can reboot into Linux, which should be able to mount the /dev/hda6 partition as before, and you can resize the *ReiserFS* filesystem 'online', that is, without unmounting the filesystem. This is done with:

```
# resize_reiserfs -f /dev/hda6
```

Without any size options, **resize\_reiserfs** will resize the filesystem to the size of the partition which the filesystem is on. **LXF**

course, you may want to check the drive with another CD, to ensure that neither the drive, nor the CD from Corel, are causing the problems you are experiencing.

## Reiser resize

**Q** I have deleted a /dev/hda3 partition with no data on it. This partition was an error in judgment that I made when I first installed SuSE 7.3 Pro on my new dual boot (Windows 2k) system some months ago. I have also removed the reference to the partition in the /etc/fstab. The result is seen below:

Device	Boot	Start	End
Blocks	Id	System	
/dev/hda1	1	2295	
18434556	c	Win95 FAT32 (LBA)	
/dev/hda2		2360	4866
20137477+	f	Win95 Ext'd (LBA)	
/dev/hda5		2360	2551
1542208+	82	Linux swap	
/dev/hda6		2552	4866
18595206	83	Linux	

Now, I would like to increase the size of my /dev/hda6 partition, which is a *ReiserFS* partition. I imagine that I would need to create a *ReiserFS* partition to replace the 500MB /dev/hda3 and somehow merge the two partitions, but I have to admit that I am clueless as to how to go about this. As I have a dual boot box, it is also possible for me to boot into Windows 2000 and use *Partition Magic*, but as far as I am aware, I can not create a *ReiserFS* partition with it.

What strategy should I use and please pencil it out. I need to learn

## 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: [lxformat@futurenet.co.uk](mailto:lxformat@futurenet.co.uk)

## Benefits – why subscribe?

- **Subscribe by Direct Debit – it's easy! Fill in the form below or call 0870 444 8645\* and have your bank details ready.**
- **Huge savings on a year's supply of *Linux Format*.**
- **Subscription copies are sent hot off the press – be one of the first to get yours!**
- **Guarantee your copy – never miss an issue.**
- **All you need to do is fill in the form. If you live in the UK, you don't even need a stamp!**

\* or call the overseas number shown at the bottom of the coupon, below

### RISK-FREE GUARANTEE!

You can cancel your subscription at any time and receive a full refund on all unmailed issues, no questions asked

Title \_\_\_\_\_ Initials \_\_\_\_\_ Surname \_\_\_\_\_  
 Address \_\_\_\_\_  
 \_\_\_\_\_ Postcode \_\_\_\_\_ Country \_\_\_\_\_  
 Email address \_\_\_\_\_  
 Daytime telephone number \_\_\_\_\_

Instruction to your Bank or Building Society to pay Direct Debits



*Media with passion*  
 Future Publishing Ltd, Cary Court,  
 Somerton, Somerset, TA11 6BR



Originator's Identification Number  
7 6 8 1 9 5

Please send this form to address shown

**1.** Name and full postal address of your Bank or Building Society branch

To: The Manager \_\_\_\_\_  
 \_\_\_\_\_ Bank/Building Society  
 Address \_\_\_\_\_  
 \_\_\_\_\_ Postcode \_\_\_\_\_

**2.** Name(s) of account holder(s) \_\_\_\_\_

**3.** Branch sort code \_\_\_\_\_  
(from the top right-hand corner of your cheque)

**4.** Bank or Building Society account number \_\_\_\_\_

**5.** Instruction to your Bank or Building Society

Please pay Future Publishing Direct Debits from the account detailed on this Instruction subject to the safeguards assured by the Direct Debit Guarantee. I understand that this instruction may remain with Future Publishing and if so, details will be passed electronically to my bank or building society.

Signature(s) \_\_\_\_\_ Date \_\_\_\_\_

Ref No (Office use only) \_\_\_\_\_

Banks and Building Societies may not accept Direct Debit Instructions for some types of account.

You may photocopy this blank form and then complete it. You cannot fax or e-mail Direct Debit forms.

QUOTE ORDER No. LXFP30

# LINUX FORMAT

Please choose your method of payment **1** or **2**

- 1.**  Direct Debit (UK only – double CD Edition) – I wish to pay £27.45 every 6 months  
 or  Direct Debit (UK only – DVD Edition) – I wish to pay £29.95 every 6 months

**Please complete the form on the left**

I understand that I will receive 13 issues during the next 12 months.

**2. Cheque/credit card 13 issues**

	CD	DVD
<b>UK</b>	<input type="checkbox"/> £54.90	<input type="checkbox"/> £59.90
<b>Europe (including Eire)</b>	<input type="checkbox"/> £78.90	<input type="checkbox"/> £83.90
<b>North America</b>	<input type="checkbox"/> £78.90	<input type="checkbox"/> £83.90
<b>Rest of the World</b>	<input type="checkbox"/> £91.90	<input type="checkbox"/> £96.90

- Cheque (payable to Linux Format. Sterling cheques drawn on a UK account.)  
 Visa  MasterCard  American Express  
 Switch Issue No/Valid date \_\_\_\_\_

Card No \_\_\_\_\_ Expires \_\_\_\_\_

Signature \_\_\_\_\_ Date \_\_\_\_\_

**UK READERS RETURN THIS COUPON TO:**

*Linux Format*, Subscriptions, Future Publishing Ltd, Freepost BS4900, Somerton, Somerset, TA11 6BR.

**OVERSEAS READERS RETURN (POSTAGE PAYABLE) TO:**

*Linux Format*, Subscriptions, Future Publishing Ltd, Cary Court, Somerton, Somerset, TA11 6TB, UK

**WEB:** [www.futurenet.com/secure/subs1.asp?ID=92](http://www.futurenet.com/secure/subs1.asp?ID=92)

Please tick this box if you would prefer not to receive information on other offers.

**ORDER HOTLINE ☎ 0870 444 8645  
OR FROM OVERSEAS ON +44 1458 271 100**

**OFFER ENDS: 14.8.02**

# EVERY ISSUE IS A SOFTWARE BONANZA!

Every issue *Linux Format* comes with either 2 CDs or a DVD crammed full of the latest Linux software. Get all this delivered direct to your door every month with a subscription costing from just £27.45 for six months!

In the last six months you would have received on DVD:

The latest **FULL LINUX DISTROS** including ● Mandrake 8.2 ● YellowDog ● Redmond Linux ● Trustix Secure Linux ● Sorcerer Linux ● Libranet ● Red Hat 7.2  
**PLUS** Over **750** of the best Linux apps and utilities.

Don't miss these great software packages – save yourself hours of download time with every issue of *Linux Format*!

« **Subscribe from just**  
**£27.45**

**every 6 months by Direct Debit**  
**saving £22.97 a year (CD issue)**

DVD issue from £29.95 for 6 months



# Coverdisc

Neil Bothwick is your guide through the wonders of this month's jam-packed Linux Format CD. Get ready to refresh your desktop.



## 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

Whether you have the bought the double CD version or the DVD copy of this issue of *Linux Format*, we have plenty on the cover discs for you. We'll take a look at the

extras that DVD users get later, first we will concentrate on what everyone gets.

When *Linux Format* came with a choice of one CD or one DVD, we tended to strip the RPM and Debian packages from the CD

version, to allow us to fit more packages on a single CD. Now that CD users have twice the space we no longer do that. Where there is an option to install from a package instead of compiling from source, CD users generally get that as well.



GNOME2 is here – almost. The final version wasn't ready when we went to press, so here's the next best thing.

## Desktop/GNOME2

We had an early beta of GNOME2 on the DVD some months ago, but it was of more interest to developers than users. GNOME2 was scheduled for release in time for this month's CD, but this was delayed slightly at the last minute, so we have the release candidate for you. This has all the release version functionality but may not be quite as stable on all systems.

Before you dive into the Desktop/GNOME2 directory, take a look in Desktop/GNOMEInstallationGuide. This has detailed instructions on installing GNOME2. You MUST read this first, installing the 2.0 files over the top of a working GNOME 1.4 setup could result in neither working properly. Some elements of GNOME 1.4 must be uninstalled first (you can always reinstall them from your distribution CDs if you want to go back).

## Graphics/QCad

CAD (Computer Aided Design) is generally considered a specialist application. While it is often used for professional design of commercial products, there are other uses too. Unless you are one of those for whom DIY means Download It Yourself, you may occasionally need to sketch out an idea or a layout. You don't want a high end package for this, just

something you can start up and use without having to learn a new skill.

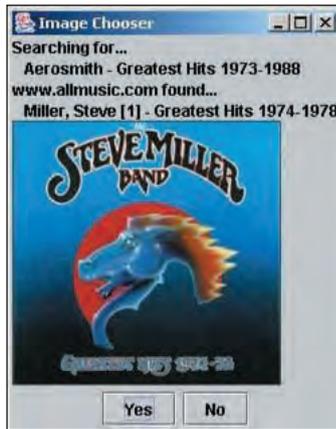
*Qcad* could be just the program you need the next time you want to redesign your kitchen or garden. It is intended to be "a stable, fast and easy to use CAD for everyone", adding "One doesn't need any knowledge of a CAD program to start working with *Qcad*". Of course, no program is ever really as easy as the its author claims, but this seems be a lot more accessible than other similar programs.

## Graphics/Transcode

One of the points raised by the video editing round up a couple of months ago was the diversity of video file formats. *Transcode* is a tool to help address this problem, by providing conversion between various formats, as well as some basic manipulation of the video streams. *Transcode* is a console program, making it ideal for batch processing files. However, expect to see various GUIs for it in the near future, as happened with the various CD and audio processing tools. There are already programs that make use of *transcode*. *DVDrip*, which does exactly what is says on the tin, uses *transcode* along with the *dvdread* and *dvdcss* libraries. *K3b*, previously an audio CD creation tool, now also handles video thanks to *transcode*, and you'll find a copy of this in the Desktop directory.

## Sound/AlbumCover Grabber

Programs like *Grip* (see *LXFDVD28*) make converting your CDs to MP3 or Ogg Vorbis files a piece of cake. As part of this process, they download track data from the freedb.org so that files and directories are named after



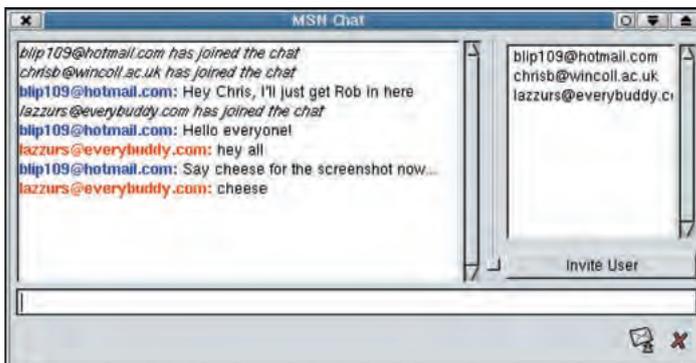
**Brighten up your MP3 collection, let *Album Grabber* add scans of the CD covers.**

the tracks and albums. *Album Grabber* takes the process a stage further, by also grabbing an image of the CD's cover from [www.allmusic.com](http://www.allmusic.com). You could do that yourself for individual CDs but this program goes a stage further and will attempt to find covers for your entire digital music collection.

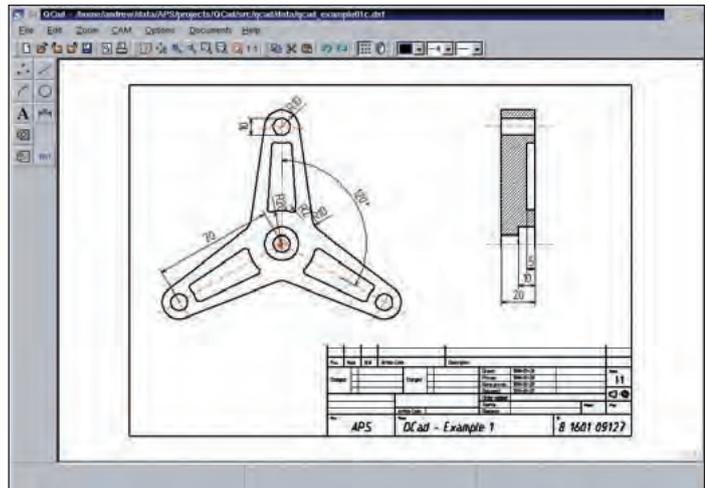
*Album Grabber* expects your music collection to be in a certain format. Your root music directory should contain artist directories, and underneath each artist directory should be album directories. This will sound familiar if you already use *Grip*, as this is the format it uses to save CD tracks. *Album Grabber* will attempt to locate the correct album cover and present it to you in a window along with 'yes' and 'no' options. The program is multi threaded, so you can let it run for awhile and then come back and make your decisions.

## Internet/Everybuddy

First there was IRC, for one to many messaging, then came ICQ for one to one chat. Then the flood gates opened



**Speak to your friends no matter which instant messaging service they use, with *Everybuddy*.**



**CAD made easy? See for yourself with *Qcad*.**



**Build your own web online store with *123tkShop*.**

and we got MSN, AOL Instant Messenger, Yahoo Messenger, Jim and Billy-Bob's Instant Chat and so on. Of course, each of your friends uses a different instant messaging service so you end up having to run many clients, or you could try an all-in-one app. *Everybuddy* is a universal instant messaging program that handles

various systems, including IM, ICQ, MSN, Yahoo! and Jabber as well as IRC. *Everybuddy* won't create accounts for you on the various services, you'll have to do that on their web sites. Once you have an account with a service, *Everybuddy* will connect to it.

Speaking of Yahoo. They are discontinuing their POP3 service. If



## Intel Compilers

Installing from your coverdisc

The Intel C, C++ and Fortran compilers are in the Magazine/Intel directory of the DVD or the first CD.

To install the C/C++ compiler type

```
tar xf /mnt/cdrom/Magazine/Intel/
l_cc_p_6.0.139.tar
./install
```

and follow the prompts.

For the Fortran compiler, type

```
tar xf /mnt/cdrom/Magazine/Intel/
l_fc_p_6.0.140.tar
./install
```

Before you can use these evaluation versions, you will need to register on Intel's web site to receive a licence key to activate the compilers. To do this, point your web browser at [www.intel.com/software/products/distributors/linuxenterprise.htm](http://www.intel.com/software/products/distributors/linuxenterprise.htm). Note that this page is unavailable at the time of writing, it is due to go live on July 1st.

If you have any problems with it, check for information at [www.linuxformat.co.uk](http://www.linuxformat.co.uk)

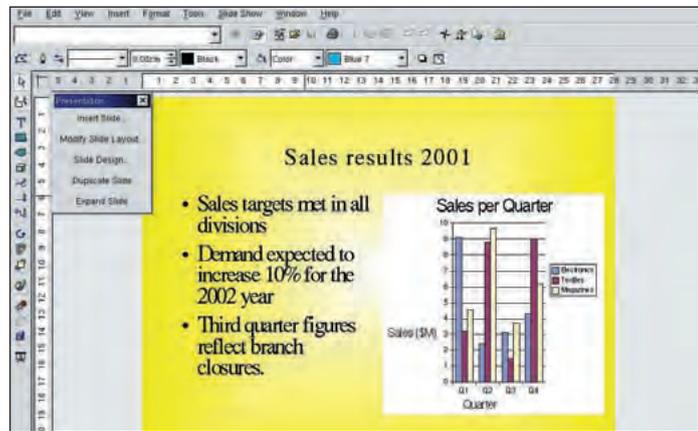
# LinuxFormatCoverdiscCD

« you have a Yahoo mail address, you'll have to use their web site to read your mail, and their adverts. At least, that's what they want, but it wasn't long before some enterprising coders came up with a solution. *YahooMailSucker* (*YoSucker* for short) grabs mail from the web site and puts it in your local mailbox. DVD users also get *FetchYahoo*, which does a similar job.

## Server/123tkShop

The Internet provides a way for businesses to reach customers they would never reach using other, more

traditional, forms of trading. If you have an existing business, already able to deal with customers and orders, adding an online facet to it is not a major task. However, you will need some e-commerce software to run on your web site, to run the catalogue, handle product searches and process orders. *123tkShop* is such a package. It is aimed at small businesses and shops and offers an easy way to add online trading facilities to a web site hosted on an ISP's server. You don't need your own server to run it, as long as you have PHP and a suitable database server available. All



SOT-Office is a full-featured office suite, partly based on OpenOffice.org.

## » CD CONTENTS AT A GLANCE

### Disc A

#### Magazine

<b>Answers</b>	Files relating to the Answers section of the magazine
<b>Emulators</b>	All the files mentioned in this month's Emulation article.
<b>Gnumeric</b>	The files for our Gnumeric tutorial
<b>HotPicks</b>	All the programs covered in this month's HotPicks section.
<b>Intel</b>	Evaluation versions of Intel Software Development Products
<b>Kylix</b>	Example files from the Kylix tutorials
<b>Officesuites</b>	Office software covered in our roundup
<b>Perl</b>	Example scripts from this month's Perl tutorial

#### Distros

<b>LRs-Linux</b>	LRs-Linux is based upon Linux From Scratch (LFS)
<b>ttyLinux</b>	Minimalist Linux distro that can fit in 4MB of disk space

#### Graphics

<b>Dia</b>	Creates diagrams of all kinds
<b>GImageView</b>	GTK+ based image viewer
<b>QCad</b>	Open-source 2D CAD system
<b>QVV</b>	Small Qt-based image viewer
<b>Transcode</b>	Text console video-stream processing tool
<b>Veeyaj</b>	A video tracking tool for X Windows and the Linux console

### Disc B

#### Desktop

<b>AcerHotkeyDriver</b>	Access the launch keys on Acer Travelmate notebooks
<b>ActiveMenu</b>	XEmacs module that saves a few extra lines of screen space
<b>APCComm</b>	Transfer files between an Amiga and a Linux PC
<b>Circulation</b>	GTK desktop application to store contact information
<b>ConfiguratorForGnome</b>	Easy editing of advanced GNOME settings
<b>Eddi</b>	Powerful and easy-to-use text editor for X
<b>GFileRunner</b>	GNOME file manager aimed to be a replacement for Nautilus
<b>GNOME2</b>	The release candidate for the forthcoming GNOME 2
<b>GNOMECommander</b>	File manager similar to Midnight Commander in text mode
<b>GNOMEInstallationGuide</b>	How to install a stable GNOME system and applications
<b>Golem</b>	Fast, lightweight, and customizable Window Manager
<b>Grutatxt</b>	Plain text to HTML converter
<b>IceWM</b>	Window manager designed for speed, usability & consistency
<b>K3b</b>	CD-burning program with similar interface to Ahead's Nero
<b>Kalculate</b>	User-friendly calculator for the desktop
<b>KNetScan</b>	Full-featured frontend for nmap, ping, traceroute and whois
<b>KNfsPlugin</b>	KDE plugin for configuring a NFS server
<b>KPortageMaster</b>	Graphical frontend for the Gentoo portage system
<b>LinTar</b>	Similar to WinZip for Windows
<b>Peeper</b>	View an area of the screen at a certain magnification
<b>phpSysInfo</b>	PHP script that parses and displays data from /proc

**Player Stage**  
**Rigel**  
**SimpleCDR-X**  
**Super-sed**  
**WMND**  
**WMnetmon**  
**WMThemer**  
**Yeasyconfig**

Networked interface to robots and their sensors  
 Personal information manager for X  
 Frontend for CD writing, mastering, and audio manipulation  
 Version of sed that supports a few new features  
 Dockapp for monitoring multiple network interfaces  
 Dockapp that monitors up to 40 hosts or services  
 Window Maker dockapp that switches between WM themes  
 PPC-only tool to help get a working XF86Config file

#### Development

<b>Ccl</b>	General purpose Java library
<b>CyberNekoDTDConverter</b>	Parses DTD files and converts the information into XML
<b>CyberNekoHTMLParser</b>	Enables Java programmers to parse HTML documents
<b>dev_conPhpBase</b>	Authentication, date, file, string, and network functions
<b>DiffXML</b>	Collection of XML diff and patch utilities
<b>Fastdep</b>	Generates dependency information for Makefiles
<b>GenDist</b>	Create your own special distribution
<b>Gnocl</b>	GTK/Gnome extension for the programming language Tcl
<b>GnosisUtils</b>	Several Python modules focused on XML processing
<b>InstantJ</b>	Compile and execute Java code on the fly
<b>IsForth</b>	x86 Forth implementation coded in pure assembly
<b>Janino</b>	Embedded Java compiler for run-time compilation purposes
<b>JcronTab</b>	Scheduler system written in Java
<b>libpng</b>	Routines to create and manipulate PNG format graphics files
<b>libtecla</b>	Provides programs with interactive command line editing
<b>NodeNet</b>	Python code for dealing with networks of connected nodes
<b>OpenCL</b>	Library of cryptographic algorithms, written in C++
<b>PyChecker</b>	Finds common bugs in Python source code
<b>PycURL</b>	Python module interface to the cURL library
<b>PyOSD</b>	Display 'On Screen Display' messages through Python
<b>RPMLint</b>	Check for common errors on RPM packages
<b>SixlegsPNGLibrary</b>	Java 1.1-compatible PNG decoder
<b>Zoinks</b>	Programmer's editor and development environment

#### Games

<b>Angband</b>	Single-player rogue-like dungeon exploration game
<b>Exult</b>	Play Ultima7 on modern operating systems
<b>Ki11egg</b>	Xiangqi (Chinese chess) program including AI and a GUI
<b>Kpictorial</b>	A logical game to reconstruct an image
<b>Papaya</b>	GTK-based MUD client for Linux, Solaris and Windows
<b>SolarWolf</b>	An action/arcade game written entirely in Python
<b>Travtrack</b>	Helps a Traveller referee keep track of essential data
<b>XKobo</b>	Fast paced, addictive, multiway scrolling shoot-em-up

#### Internet

<b>AlvarosMessenger</b>	Tcl client for MSN Messenger
<b>Animail</b>	Multiserver POP3/APOP/IMAP4Rev1 mail retrieval utility
<b>BKFTP</b>	Easy access to FTP with a minimalistic graphic UI
<b>Circle</b>	Peer-to-peer file sharing, instant messaging and news
<b>Cloish</b>	Make your Web surfing anonymous
<b>DConnectDaemon</b>	Daemon that acts as a Direct Connect hub

configuration and administration is done with a web interface, including backups of the database tables.

Various themes are available, so you don't have to worry about your site looking exactly the same as every other site using *123tkShop*. Being written in PHP, if you have some knowledge of the language you can easily adapt it to suit your needs, although this is a nicety rather than a necessity. This is a fairly early version, so it doesn't offer all of the planned features, but it should be enough to help you get started in online trading.

## Server/Portmon

One of the good things about Linux is that if a process crashes it doesn't generally affect the rest of the system. However, there is a downside to this. If a server you are running terminates unexpectedly, you may not realise anything is wrong until someone complains that they cannot connect.

This is where a program like *Portmon* can help. *Portmon* sits in the background, connecting to a list of servers and services specified in its configuration file. By default it runs through this list once per hour

although this can be changed. When a server fails to respond, one of a number of actions can be taken, sending you an email or ICQ message, popping up a window, beeping or playing an annoying sample.

*Portmon* is not limited to checking the machine it is running on, there wouldn't be much point in it sending you an email to warn you the mail server had failed, it will happily keep and eye on a range of servers running on a network of machines. No software is required on the machines being monitored, they don't even have to be running Linux.

## Office Suites

Unless you are one of those people who read magazines from the back forwards, you cannot have failed to see our mammoth feature on office software. All of the open source software covered in here is on the CD, in the Magazine/Officesuites directory. There is also another suite of office software on the CD. *SOT-Office* is partly based on *OpenOffice.org*, but that doesn't mean it's the same, in the same way that browsers based on *Mozilla* can be quite different to use. *SOT-Office* is in the Office directory. 

<b>DConnectGUI</b>	DC-GUI is a QT GUI Direct Connect filesharing client.
<b>Everybuddy</b>	Universal Instant Messaging client
<b>Galeon</b>	GNOME Web browser based on the Mozilla rendering engine
<b>InitScriptsVPNExtensions</b>	Easily implement VPNs using PPP over SSH or PPTP
<b>JFtp</b>	Graphical FTP client written in Java
<b>Lftp</b>	Sophisticated command line based FTP client
<b>Mozilla</b>	The final release, and the next alpha version
<b>Onlinecalc</b>	Summarises your online time, transfer volume etc.
<b>OpenVPN</b>	Robust and highly configurable VPN daemon
<b>psyBNC</b>	Multi-user, permanent IRC-Bouncer with many features
<b>SDSC_GTSecureFTP</b>	Make a secure connection to an FTP daemon via SSL
<b>Signit</b>	Customisable signature creator
<b>SpamScan</b>	Filters email on content using keywords set by the user
<b>WMget</b>	Monitor or control multiple downloads from within a dockapp
<b>WMpop</b>	Window Maker dock app that monitors a local or POP3 mailbox
<b>YahooMailSucker</b>	Download Yahoo Mail messages to your local inbox
<b>Yeemp</b>	Decentralized instant messaging system
<b>Zebedee</b>	Establish an encrypted, compressed tunnel for data transfer

## Office

<b>AngleMailForPhpGroupWare</b>	An alternative email module for phpGroupWare
<b>ePiX</b>	Powerful, flexible, lightweight, text-based preprocessor
<b>MimerDesk</b>	Groupware environment designed for a wide variety of uses
<b>OpenSymphonyWorkflow</b>	Uses range from document management to order processing
<b>SOT-Office</b>	Office suite, fully compatible with other major office apps
<b>Taglog</b>	Keep track of time spent on projects

## Server

<b>123tkShop</b>	e-Business framework to let small companies set up webshops
<b>AdaptiveWebsiteFramework</b>	Create dynamic personalised Web sites
<b>Bounceweb</b>	Approve bounced Majordomo messages through a Web interface
<b>CGIForum</b>	Highly configurable template-based discussion board
<b>ChatEverywhere</b>	Build a powerful but easy-to-use chat area on a Web site
<b>DCVS</b>	Perl/MySQL CVS server supporting the pserver protocol
<b>Dnsmasq</b>	Lightweight DNS forwarder for small networks
<b>DNSTools</b>	Browser-based DNS configuration and administration utility
<b>Drupal</b>	Full-featured content management/discussion engine
<b>Gween</b>	Something like a Mason port to PHP
<b>HitCount</b>	CGI hit counter
<b>ibWebAdmin</b>	Web based admin for InterBase/Firebird database servers
<b>Informium</b>	Advanced news script for use within a Web site
<b>Lurker</b>	Mailing list archiver with Google-style searching
<b>MHonArc</b>	HTML mail archiving with index, mail thread linking etc.
<b>MiddleMan</b>	Proxy server to remove unwanted content & increase privacy
<b>MIMEDefang</b>	MIME email scanner to protect Windows clients from viruses
<b>MIMETypes</b>	Generates mime.types files for Apache
<b>mmFTPd</b>	Secure FTP server that runs as a normal user

<b>MySQL</b>	Widely used and fast SQL database server
<b>NotFTP</b>	Fully-featured and easy-to-use Web-based HTTP-FTP gateway
<b>OASIS</b>	High-performance ad server built using PHP on top of MySQL
<b>OpenDCHub</b>	Hub software for the Direct Connect network
<b>phpMySearch</b>	Provides a search feature for your own Web site
<b>pIRcd</b>	One of the few IRC daemons with SSL support
<b>Portmon</b>	Monitor network services and warn when a server fails
<b>ProFTPD</b>	Proven, high-performance, scalable FTP server
<b>pyWM</b>	Webmail software written in Python
<b>Sendmail</b>	The most used Mail Transfer Agent
<b>Spyce</b>	Apache plugin for Python-based dynamic HTML scripting
<b>Stupid-FTPd</b>	FTP server with normal functionality and a commandline mode
<b>Twisted</b>	Event-based framework for Internet applications
<b>VipulsRazor</b>	Spam detection and filtering network
<b>W-Agora</b>	Database-driven web-base communications system

## Sound

<b>AlbumCoverGrabber</b>	Grabs album covers from www.allmusic.com
<b>Audacity</b>	Cross-platform multitrack audio editor
<b>c3Dse</b>	Enable the 3D Stereo Enhancement of your SoundBlaster
<b>DivineIntervention</b>	Music tempo measurement and adjustment
<b>EasyTAG</b>	View and edit tags for MP3, MP2, FLAC and Ogg files
<b>EmptytreeSeedy</b>	Distributed digital music system
<b>GnomeMP3</b>	Rename and tag MP3 files
<b>GSMP</b>	The General Sound Manipulation Program
<b>KMusicChoice</b>	Fetches information about songs playing on MusicChoice
<b>LAME</b>	LAME Ain't an MP3 Encoder... but it is.
<b>MyMP3s</b>	Controls the playing of MP3s on a server
<b>SoundTracker</b>	Pattern-oriented music editor
<b>SQMixer</b>	Audio mixer for OSS-compatible drivers that uses QT
<b>Trackbox</b>	Trackbox is a pure music server
<b>Umix</b>	Adjust volumes and other features in soundcard mixers
<b>ZoltanPlayer</b>	ZoltanPlayer is a music playing daemon

## System

<b>AIDE</b>	Advanced Intrusion Detection Environment
<b>Auto-autofs</b>	Searches devices and partitions to make an automounter map
<b>Cheops-ng</b>	Network tool for mapping and monitoring your network
<b>CoherentPrintingSystem</b>	Foomatic/GhostScript-based print spooling system
<b>CUPS</b>	Portable printing layer for UNIX(r)-based operating systems
<b>Fireflifer</b>	Firewall tool built on top of the iptables framework
<b>iBackup</b>	Automated backup of system configurations
<b>LILO</b>	Boot loader for Linux/x86 and other PC operating systems
<b>LinuxNTFS</b>	Drivers for the filesystem used by Windows NT, 2000 and XP
<b>mkAutoSMB</b>	Autodiscovers the local Windows network using smbclient
<b>MonitorIT!</b>	A multi-purpose monitoring tool.
<b>NewSyslog</b>	Manages the rotation and archiving of log files
<b>Sky-Fire</b>	Secure, easy-to-use Linux firewall
<b>Sonar</b>	Network reconnaissance utility
<b>TheEDDIETool</b>	System and network monitoring and performance analysis tool
<b>Xinetd</b>	Replacement for inetd, the internet services daemon

# Coverdisc

**Neil Bothwick** is your guide through the wonders of this month's jam-packed *Linux Format DVD*, and its comprehensive Perl archive.



The last few months have seen updates to some major distributions or desktops. This has resulted in as much as half of the DVD being given over to a single product. That's good, it's one of the reasons for providing a cover DVD, the ability to include such large projects. However, this month there is no such elephantine release, so we have taken a slightly different approach to the DVD.

The largest item on this month's DVD is not a single product but an archive. One of the web's most useful resources for programmers and users alike is mirrored on this month's DVD, bringing a speed and convenience to

everyone that surpasses even a T1 Internet connection.

## Development/CPAN

Perl is an acronym, the "P" standing for "Practical". It could equally well be "Pervasive". Perl gets everywhere, so many useful programs are written in it, from small utilities to full blown distribution installers. One of the things that makes it so useful is that almost every problem has been tackled before, someone has come up with a solution, and often released it as a module. This is good for programmers because they have a vast library of existing code to draw on, but it can be a problem for users. You can guarantee that the really useful script you've just found, that does exactly what you need, requires a module not installed on your system. This means trekking over to CPAN (the Comprehensive Perl Archive Network) to search for and download the modules you need. The download time for most modules is minimal, you'll spend more time dialling up and navigating the site... until now.

Now you only need to put your Linux Format DVD in the drive, because we have a complete mirror of CPAN on this month's disc. As a user, this means you no longer have to search online for any modules you need. On the other hand, developers can browse the archive to find a wheel that fits their needs instead of reinventing their own. CPAN is regularly updated, and it's a huge archive. We can't put it on the DVD every month as we do with the Linux Documentation Project, but we intend to put updates on the DVD in future. If you have the hard drive space to store the CPAN archive, you'll be able to copy in the updates to give yourself a complete, up to date (well, as up to date as magazine publishing schedules allow) CPAN mirror.



Whether you're writing Perl scripts or trying to run them, you'll find all the modules you could possibly need in our CPAN archive.

## Distros/GentooLinux

The "major" distributions, like Red Hat, Mandrake and SuSE, put a great deal of work into producing a selection of packages that will cover almost every need, combined with an installation and subsequent configuration that is as easy as possible. This is great, it means Linux is more accessible than ever, but it's not what everyone wants (this isn't a bad thing when you

consider that the closest thing to what everyone wants is apparently Windows XP!) Gentoo Linux is no relation to the gentoo file manager. While there is no link between them, they do have something in common besides the name. They both aim to provide an efficient way of working for those that are willing to make the effort to set them up.

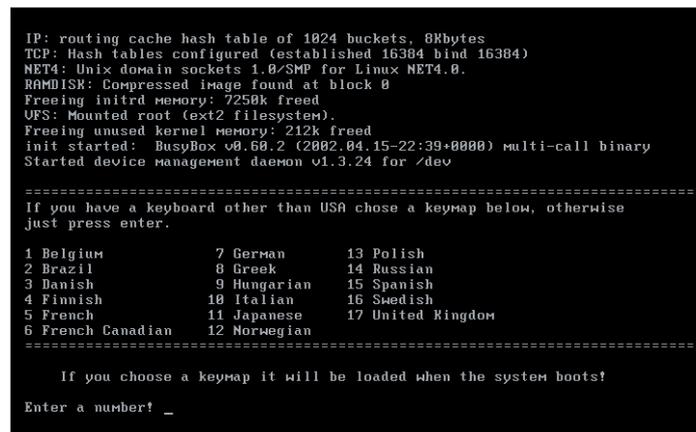
Gentoo doesn't provide a hand holding GUI for installation, not even

On the DVD  
DVD  
COM

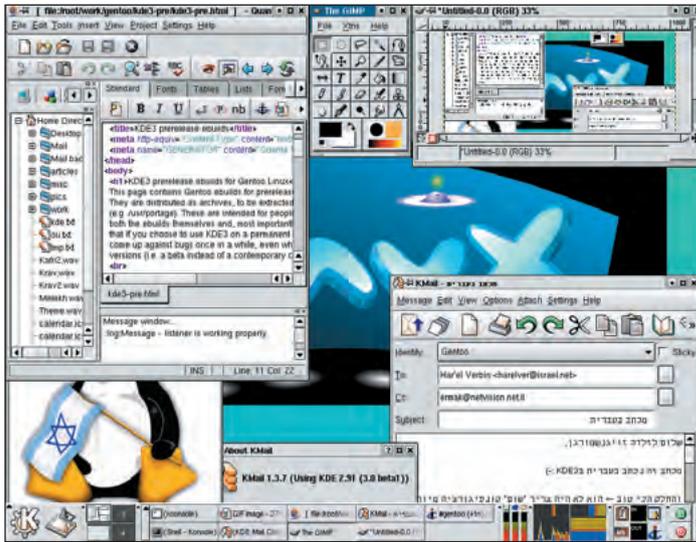
Wherever you see this logo it means there's related stuff on the DVD

## Important notice

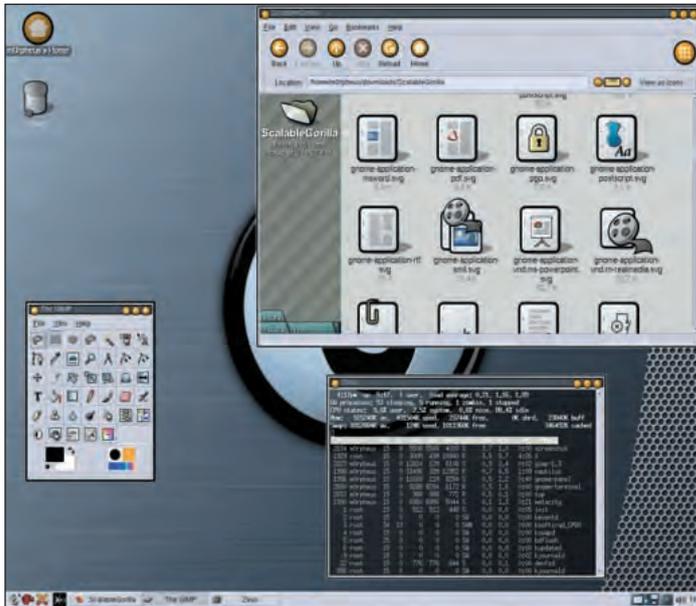
**Before you even put the CD in your drive, please make sure you read, understand and agree to the following: The Linux Format CD is thoroughly tested for all known viruses, and is independently certified virus-free before duplication. We recommend that you always run a reliable and up-to-date virus-checker 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.**



It's not pretty, but installing Gentoo Linux is reasonably straightforward if you follow the instructions.



The installer may now be much to look at, but a fast and efficient distribution doesn't have to look boring. Below: Gentoo in all its glory.



an ACSII-based one. That said, it's not particularly difficult to install for anyone with a reasonable knowledge of Linux, provided you take the time to read the documentation. This is NOT a distribution for someone who just wants to boot from a disc and click a few buttons, it provides control for the user, but that control comes at the price of having to work a bit harder. The other benefit of installing a distribution like Gentoo is that it helps you learn more about how a Linux distribution operates, even if you don't install it as your main operating system.

After booting from the CD and running through the keyboard selection, PCI hardware detection and installation of a bare bones console system, you are told to read the install.txt file. This gives detailed instructions on how to proceed, I found it best to switch to another console (**Alt-F2**) so I could read the instructions in one and follow them in the other. Then it's simply a matter of working through the directions.

After setting up the networking and filesystems, it's time to install the rest. There are three "stages" you can install, stages 1 and 2 are minimal installs so the rest can be done from the network or from source code. Stage 3 provides a full working environment and is probably the one most people will want to use first time. This is effectively a pre-built environment, installation consists of *cd'ing* to your new root partition and unpacking the tarball. You are almost finished now. The only major task left is compiling a kernel (see, I told you it wasn't a simple point and click install). As with all the previous stages, this is clearly documented in install.txt, follow the instructions and you should be

fine. Then you edit a couple of configuration files, set up *GRUB* (all clearly documented) and you're ready to reboot.

## Games/Boson

Perl modules, alternative Linux distributions, it all sounds a bit serious. What about some fun? Don't worry, we have games on the DVD too. *Boson* is a real-time strategy game, like *Command & Conquer* or *StarCraft*. It is built on top of the *libkdegames*, KDE and Qt libraries. You don't need to be running KDE to play the game, but it must be installed. This is a two player game, so you'll need a friend to be your enemy. There is a slight problem with the tarball archive which may cause problems with the standard installation procedure. The fix is simple, run **make distclean** before **./configure** and all should be well.

One game, even a good one, may not keep you entertained for long, so we have several others on the DVD. Tank Command is a fast paced 2D shoot-em-up with a strong strategic element.

## Internet/Mozilla

*Mozilla 1.0* has finally arrived, on both the CD and the DVD. But if you thought that was the end of it, think again. The first full release is not the end, merely the end of the beginning. DVD users get 1.0 in both source and binary packages, but we also have the first public version of *Mozilla 1.1* on the DVD.

## Sound/Normalize

If you have a large collection of music in MP3 or Ogg Vorbis files, you have



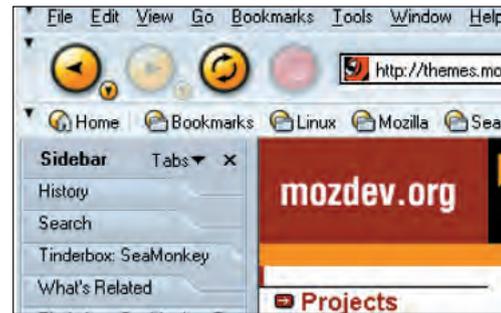
It's not all serious stuff on the DVD, There are plenty of games to play, including the real-time strategy *Boson*.

# LinuxFormatCoverdiscDVD

probably noticed that the volume can vary quite a lot between files. This is particularly noticeable when you turn up the volume for a quiet file and it's followed by a loud one. *Normalize* is a solution to this. It processes a number of audio files and adjusts the volumes to match. It does this in one of two ways. MP3 files can have a relative volume tag in the MP3 header, although this isn't recognised by all

players. The more universal method involves adjusting the volume of WAVE files. This means MP3 and Ogg files must first be uncompressed before converting and recompressing, although you can also change the format with this process. The most effective use is when ripping CDs in the first place as it can work on the raw, and highest quality, audio data from the CD. [LXF](#)

**The release version of Mozilla 1.0, the first alpha of 1.1 and a bunch of themes too, browser heaven!**



## »» DVD CONTENTS AT A GLANCE

### Magazine

Answers	Files relating to the Answers section of the magazine
Emulators	All the files mentioned in this month's Emulation article.
Gnumeric	The files for our Gnumeric tutorial
HotPicks	All the programs covered in this month's HotPicks section.
Intel	Evaluation versions of Intel Software Development Products
Kylix	Example files from the Kylix tutorials
Officesuites	Office software covered in our roundup
Perl	Example scripts from this month's Perl tutorial

### Desktop

AcrobatReader	View, navigate, and print PDF files
Autorun-KDE	Enable auto-insert notification for CDROM/CDR/DVD drives
BBUNAME	Blackbox/Fluxbox tool to display system information
Ghostsript	Processor for PostScript and PDF files
GNOMEGhostView	A frontend for Ghostscript
Ispell	Interactive spell-checking program
Kbarcode	Creating, handle, manage and print barcodes
KMouseTool	Clicks the mouse for you, reducing RSI pain
Knowde	Knowledge management tool
Ksplash	Configurable splash screen for KDE3
LingoTeach	Simple language teaching tool
mcrypt	Encrypts files or streams
Minkowsky	Calendar, address book and task management software
QTLDAP	QTLDAP is a LDAP client/browser made with Qt.
TaskJuggler	Powerful project management tool for a wide range of tasks
Themes	A selection of desktop themes and wallpapers
Twin	Turns a text terminal into a X11-style display
VisualOS	Educational visual simulator of an operating system
Xpdf	Viewer for Portable Document Format (PDF) files

### Development

AdvancedBashScriptingGuide	Comprehensive reference and tutorial on shell scripting
CPAN	The Comprehensive Perl Archive Network (CPAN)
CyrusSASL	Library to integrate secure network authentication
IMDbPHPClass	Functions to get info from the Internet Movie Database
KDevelop	C/C++ development environment
TurboToGccPortingTool	Port MS-DOS based Borland Turbo C code to GNU gcc

### Distros

DMZS-Biatchux	Portable, bootable CD-ROM-based distribution
GentooLinux	Versatile and fast Linux distribution
Securepoint	High-end and cost-effective firewall and VPN solution

### Games

Boson	Real-time strategy game like Command & Conquer or StarCraft
Freeciv	Reimplementation of the famous game of Civilization
LiquidWar	A truly original multiplayer wargame
Rhino	Othello/reversi game for the Gnome environment
TankCommand	A fast paced 2D arcade game

### Graphics

FAME	Real-time MPEG encoding
------	-------------------------

SWFTools  
Trailer  
XDrawChem

Collection of code for handling Flash .SWF files  
Create thumbnailed albums arranged in chronological order  
A program for drawing chemical structures

### Internet

DownloaderForX	Downloads files from the Internet via both HTTP and FTP
FetchYahoo	Perl script to download mail from a Yahoo! webmail account
Kopete	Flexible and extendable instant messaging system
MailArmor	Filters spam before it leaves the mail server
MozillaThemes	Themes for use with Mozilla 1.0
NetscapeCommunicator	All-in-one web browser and communications suite
Sylpheed	GTK+ based, lightweight, and fast email client
WebCleaner	A filtering HTTP proxy

### Office

CK-Ledger	Double-entry ledger accounting system for PHPGroupWare
CRMCommercialTrackingTool	Log and track customer requirements, complaints etc.
Scrbus	Easy-to-use tool for simple Desktop Publishing
WebInvoice	Website project to manage clients, billing and payments

### Server

ApacheSearchEngineLog	Extracts the search engine terms that led to one's site
COMUPrivacyGuard	Web-based shell for GNU Privacy Guard
DomainTechnologieControl	Manage various aspects of domains and subdomains
gIFtpD	FTP Daemon that is easy to setup and use
GProFTPD	GNOME frontend for the ProFTPD standalone server
Mailto.php	Protects mailto: links from address harvesters
mmailSuite	SMTP and POP3 daemons using MySQL
phpShare	A file browser written in PHP4
Ringlink	Run one or more rings of Web sites
SMPTRelayChecker	Fully configurable, multithreaded open mail relay scanner

### Sound

EarCandy	A non-Web based music jukebox
GNUsound	Sound editor supporting multiple tracks and outputs
IceSoundManager	Manage sound events, sound themes and the IceSound server
LongPlayer	Player for people with large MP3 or Ogg Vorbis collections
Normalize	Adjusts the volume of audio files to a standard level
WMmixer-DockApp	Small dockable mixer application

### System

FirestormNIDS	High performance network intrusion detection system
Fsbackup	Incremental backup creation utility
Monit	Monitor and manage daemons or similar programs
NetworkProbe	Network monitor and protocol analyzer
Srcrser	Package manager that install packages from source tarballs
UserFriendlyIPTables	Generate optimized iptables packet filter rules
WebminUsermonitor	Webmin module to view all user connections to the system

### Essentials

LDP	A complete mirror of the Linux Documentation Project, containing HOWTOs, man pages and complete books!
-----	--

# 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.

## 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-1.0.i386.rpm – This is probably a binary rpm, designed to run on x86 systems.

Someap-1.0.i386.deb – The same, but a debian package.

Someap-1.0.targz – This is usually source code.

Someap-1.0.tgz – Same as the above, tgz is abbreviated form of targz

Someap-1.0.tarbz2 – Same, but uses bzip2 compression instead of zip

Someap-1.0.src.rpm – This is also source code, but supplied as an rpm to make it easier to install

Someap-1.0.i386.RH7.RPM – A binary, x86 RPM designed specifically for Red Hat Linux

Someap-1.0.ppc.Suse7.rpm – A binary RPM designed specifically for SuSE7.x PPC Linux.

Someap-devel-1.0.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/progname-2.1.0.tgz
tar xjvf /mnt/cdrom/Desktop/progname/progname-2.1.0.tar.bz2
```

Use the first for Gzipped files, those ending in .targz 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**

## 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.

# Help wanted

Summer's here, so it's probably raining outside. Even if it is sunny, **Richard Smedley** suggests that you stay in and further the cause of Free Software.

**E**ducation seems to be a popular field for Free Software projects, which speaks volumes for the enlightened attitudes of many Free Software coders – looking to help and educate the next generation.

Tux4Kids is an umbrella project currently covering the *Tux Typing*, *Tux of Math Command* and *TuxReader* projects.

*Tux Typing* is the flagship title for Tux4Kids. It is an educational typing tutor for children. This stable and mature project runs under Linux, BeOS, and Windows, and is well worth a look at – but it is the other two projects that need more help.

*Tux of Math Command* is a math problem game at the early stages of development – see

[www.geekcomix.com/dm/tuxmath](http://www.geekcomix.com/dm/tuxmath)  
The game requires SDL.

The imagery – with ruined cities

being defended by Tux, against falling comets (containing maths problems) – has caused the developers some soul-searching since the attack on New York last September. Concerned that their images no longer seem so innocent, they have put up a poll on possible changes at:

[www.geekcomix.com/dm/tuxmath/poll.shtml](http://www.geekcomix.com/dm/tuxmath/poll.shtml)

## TuxReader

Another SDL app, *TuxReader* uses the *Festival* speech synthesis system to help teach younger children how to read. It is at a very early stage, but you can check out the code from Savannah: <http://savannah.gnu.org/cgi-bin/viewcvs/tuxread/>

[www.geekcomix.com/dm/tuxread](http://www.geekcomix.com/dm/tuxread)

All of the above projects should soon be moving to new pages hosted at [www.tux4kids.org](http://www.tux4kids.org). If you have a project that you feel should be placed



The future of *Tux of Math Command's* ruined towers is now the subject of an online poll...

under the Tux4Kids sponsorship umbrella, you can find more details on

joining the project at [www.tux4kids.org/tux4kids/dev](http://www.tux4kids.org/tux4kids/dev)

## Rosegarden

Powerful sequencer and musical editor

From **Richard Bown** and **Guillaume Laurent**, at Rosegarden:

We're always looking for people willing to participate in the project. If you're interested, here are a few clues about how to start.

### First steps

Subscribe to the *rosegarden-devel* mailing list. <http://lists.sourceforge.net/lists/listinfo/rosegarden-devel>

As with any software project, the docs are a little thin on the ground but there are some design docs and HOWTOs in the docs/ area of the source tree. You can also browse the code reference docs [http://rosegarden.sourceforge.net/code\\_doc/](http://rosegarden.sourceforge.net/code_doc/) which we try to keep up to date. If you're stuck somewhere or don't understand something then please ask on the mailing list.

<http://lists.sourceforge.net/lists/listinfo/rosegarden-devel>

### What next?

Pick something that you want to do. Make it something you're interested in and think you can help with. This is something you'll have to do for yourself, as we find being prescriptive with tasks just turns people off. Best if you choose your own job and therefore find your own motivation. (But be practical. All the time spent working on one thing is time that can't be spent working on another. If your goal is to make software that other people can use and enjoy, work on the things that will give you the best return in usability on your investment of time. Working on something concrete based on existing code is a better use of your time than trying to change large parts



Rosegarden – MIDI and audio sequencer, and musical notation editor.

of the design to support something fantastic that you dream of doing one day.) If you want to see what we're working on now then have a look at our task lists. They might be a little out

of date but they'll give you a rough idea of what we're up to at the moment. [http://cvs.sourceforge.net/cgi-bin/viewcvs.cgi/rosegarden/docs/task\\_lists/](http://cvs.sourceforge.net/cgi-bin/viewcvs.cgi/rosegarden/docs/task_lists/)



# 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)

## 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)

## 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

## 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

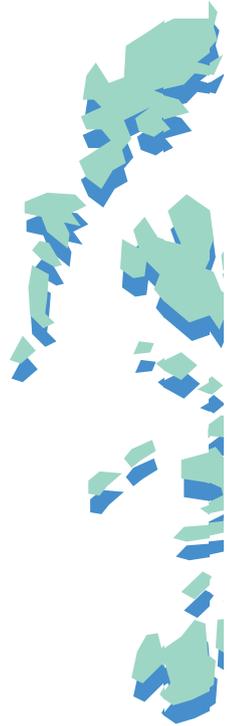
URL [www.nottingham.lug.org.uk](http://www.nottingham.lug.org.uk)  
Contact Godfrey Nix

## 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)

**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)

**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 Francis Barton

REVISED  
DETAILS**63 West of Scotland**

Contact Ray Gillian  
 Email [raygillan@fire-bug.co.uk](mailto:raygillan@fire-bug.co.uk)

NEW  
DETAILS**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.bordnet.net](http://www.linux.bordnet.net)  
 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@coventry.lug.org.uk](mailto:info@coventry.lug.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)

**All the following are  
 new details:**

**75 East Yorkshire**

Contact Daniel Gallacher  
 Email [sharkonline@whatemail.com](mailto:sharkonline@whatemail.com)

**76 Cleveland Open Source Group**

Contact Haniff Din  
 Email [openlug@digitalmedia.co.uk](mailto:openlug@digitalmedia.co.uk)

**77 Beverley**

Contact Vladimir Lukyanov  
 Email [vladimir\\_lukyanov@hotmail.com](mailto:vladimir_lukyanov@hotmail.com)

**78 Dundee & Tayside**

URL [www.dundee.lug.org.uk](http://www.dundee.lug.org.uk)  
 Contact Duncan Gauld

**79 Sussex**

URL [www.phpworld.co.uk/~swlug](http://www.phpworld.co.uk/~swlug)  
 Contact Gareth Ablett

# LinuxUserGroups

## LUG OF THE MONTH!

### Fuengirola

Although Linux is very well established in Spain, there is a language barrier for the many resident foreigners. A typical manual is hard enough to understand in English.

FuGLUG started in late 1999 as an offshoot of a Windows UG based in Fuengirola in Andalucía. The original four members struggled as complete novices to a Unix style OS. We are now an eight-member group and have

three experienced Linux and BSD users. As a result of this increase, our fortnightly meetings are much more interesting with significant exchanges of knowledge. We have an online discussion group and often meet socially.

Although we are enthusiastic Linux users, mostly we are hobbyists but do have members with professional experience of web design, video editing, programming and hardware supply.

We are English, Dutch and American, aged from under 40 to over 80. We have high hopes that the recent Linux Install Day will increase our membership, and are still looking for our first Linux chick.

Should any vacationing Linux Guru happen to drop in, we have a standing offer of free cerveza and paella.

[www.fuengirola.lug.org.uk](http://www.fuengirola.lug.org.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)

#### ALICE SPRINGS

**URL** [www.aslug.org.au](http://www.aslug.org.au)

#### MELBOURNE, VICTORIA

**URL** [www.luv.asn.au](http://www.luv.asn.au)  
**Contact** [luv-committee@luvasn.au](mailto:luv-committee@luvasn.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.ditu.org](http://www.ditu.org)

**Email** [glossary@ditu.org](mailto:glossary@ditu.org)

#### GOTHENBURG

<http://nain.oso.chalmers.se/LUG/>

**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)

#### ALBERTA

**URL** <http://calgary.linux.ca/>

#### BATON ROUGE

**URL** [www.brbug.net](http://www.brbug.net)

**Email** [dpuryear@usa.net](mailto:dpuryear@usa.net)

#### BAY AREA

**URL** [www.balug.org](http://www.balug.org)

**Email** [aftyde@balug.org](mailto:aftyde@balug.org)

#### CLARKSVILLE, TN

**URL** [www.cllug.org](http://www.cllug.org)

**Email** [tux@cclug.org](mailto:tux@cclug.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.vtlug.org](http://www.vtlug.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)

## LUGs – your local help

Welcome to our expanded User Group pages. With nearly 80 LUGs active around the UK, and several hundred more around the world, we are happy to give them more space.

LUGs are often the focal point for Free Software involvement in a community, organising events for local businesses, helping schools and even providing a good excuse to go down the pub and meet some friends.

### LUGs need you!

However the one thing a LUG needs to thrive is *you*. A LUG is a group of GNU/Linux users – no more, no less. The success of each and every LUG is down to the enthusiasm of its members. It needn't involve too much effort – just answering the odd question on the LUG mailing list, or helping to keep the website up to date – but each and every member counts.

If you want to support Free Software use in your local area then joining your local LUG is a great start. If you live in the UK have a look at [www.lug.org.uk/lugs/index.html](http://www.lug.org.uk/lugs/index.html) If not take a look at [www.linux.org/groups](http://www.linux.org/groups) We will try and keep our map (see page 111) as up to date as possible, but there is always the chance that a new LUG has appeared after we have gone to press.

As well as helping your fellow Linux users, the chances are that you will benefit from the different experiences of your fellow members.

What are you waiting for? Join in!

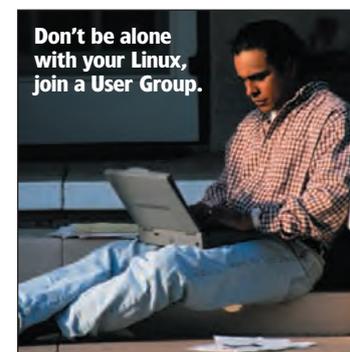
### LUG events

This spring's Install Days showed the potential both for LUG events and company involvement. LUGs can promote Linux in a variety of ways. Sheffield LUG kicked off this year with their huge Linux seminar involving IBM and SuSE, along with local businesses and a speaker from the Free Software Foundation Europe. If you run a business, then why not consider sponsoring an event?

### What, no LUG?

If there is no LUG in your area, then there is only one thing for it – start your own. As with all things Linux there is a HOWTO guide at [www.linuxdoc.org/HOWTO/User-Group-HOWTO.html](http://www.linuxdoc.org/HOWTO/User-Group-HOWTO.html) As well as a short guide at [www.lug.org.uk/lugmasters/howto.html](http://www.lug.org.uk/lugmasters/howto.html)

Set yourself realistic targets, and try and share out the workload – this will keep it a pleasure, not a burden. **LXF**



## 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)

# NEXT MONTH

## Issue 31 on sale Wednesday 14 August

### 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, David Cartwright, Andy Channele, David Coulson, Hoyt Duff, Simon Goodwin, Paul Hudson, Richard Ibbotson, Intel, Neil Lucock, Maurice Kelly, Brian Long, Charlie Stross

### COVER CD PRODUCTION

**CD Editor** Neil Bothwick  
**Senior CD Editor** Kate Hadley

### ART CONTRIBUTORS

**Photography** Dave Caudery, Getty Images, Corbis  
**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

# HURRAH FOR



Or should they be renaming it Lollywood?

Over the last few years, Linux has been making steady progress at infiltrating the CGI market, with more and more production houses won over by its reliability and performance. In a special feature we'll look not just at the high profile successes, but the underlying shift to Linux, the reasons and implications.

## Security

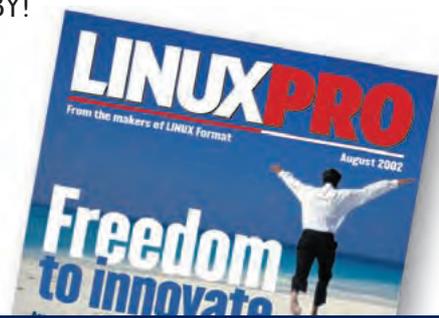
In a prelude to our forthcoming firewall extravaganza, we'll be exploring general security issues, the tools you need and more importantly, giving you the insider info on how to manage your systems securely.

## PLUS:

Reviews of Zend Studio, Quasar Accounts, Slackware 8.1, and a roundup of mailservers; tutorials on Unix text tools and CD writing, and more on Gnumeric, Intel compilers and PHP; plus RUBY!

## Including LINUXPRO

United Linux! Is a unified Linux desirable or even possible? Could it work this time, when so many 'unified' unices have failed in the past? We'll find out and tell you here!



**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

August 2002

## Freedom to innovate

Are software patents a valuable safeguard to intellectual property, or will they stifle business and open source creativity?

### PLUS:

**Experiences of a Cybercafé** – Edinburgh's Electronic Cottage investigated

**Memory matters** – techniques for programming in embedded environments

**Stand by to repel boarders** – Securing servers with PortSentry

## Contacts

### 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)

#### Photography

Superstock, Getty Images

#### Publisher

Dave Taylor

### ADVERTISING

#### Group advertising manager

Simon Moss

#### Advertising manager

Matt Stanley  
[matt.stanley@futurenet.co.uk](mailto:matt.stanley@futurenet.co.uk)

#### Deputy advertising manager

Emma Mimmack  
[emmamimmack@futurenet.co.uk](mailto: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

Yes, our Linux Pro supplement is proving so popular that you're getting another copy, and one with a few interesting new twists. We're aware that you all enjoy good case studies, and to fit our mini-mag format, we have a mini case study on a cybercafé that runs Linux. But this isn't just a standard web access coffeeshop as you'll see.

Following on from our article last month on the legal issues of Free Software licences, this month the lawyers better take cover as we try and sort out the winners and losers in the software patent lottery, with a little help from Richard Stallman.

Our sponsors this issue, The Positive Internet company, have submitted their own personal take on how Free Software has changed the way business works, we have an informative piece on coding for embedded applications and to round



things off, there's another security mini-tutorial, this time on setting up *PortSentry*. As ever, we hope you enjoy it – please feel free to email me!

**Nick Veitch** Editor  
[nick.veitch@futurenet.co.uk](mailto:nick.veitch@futurenet.co.uk)

## Contents

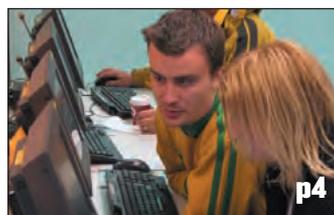
**ELECTRONIC COTTAGE:**  
 Mocha and Mozilla **p4**

**MEMORY MATTERS:** Good ideas for uCLinux coding **p8**

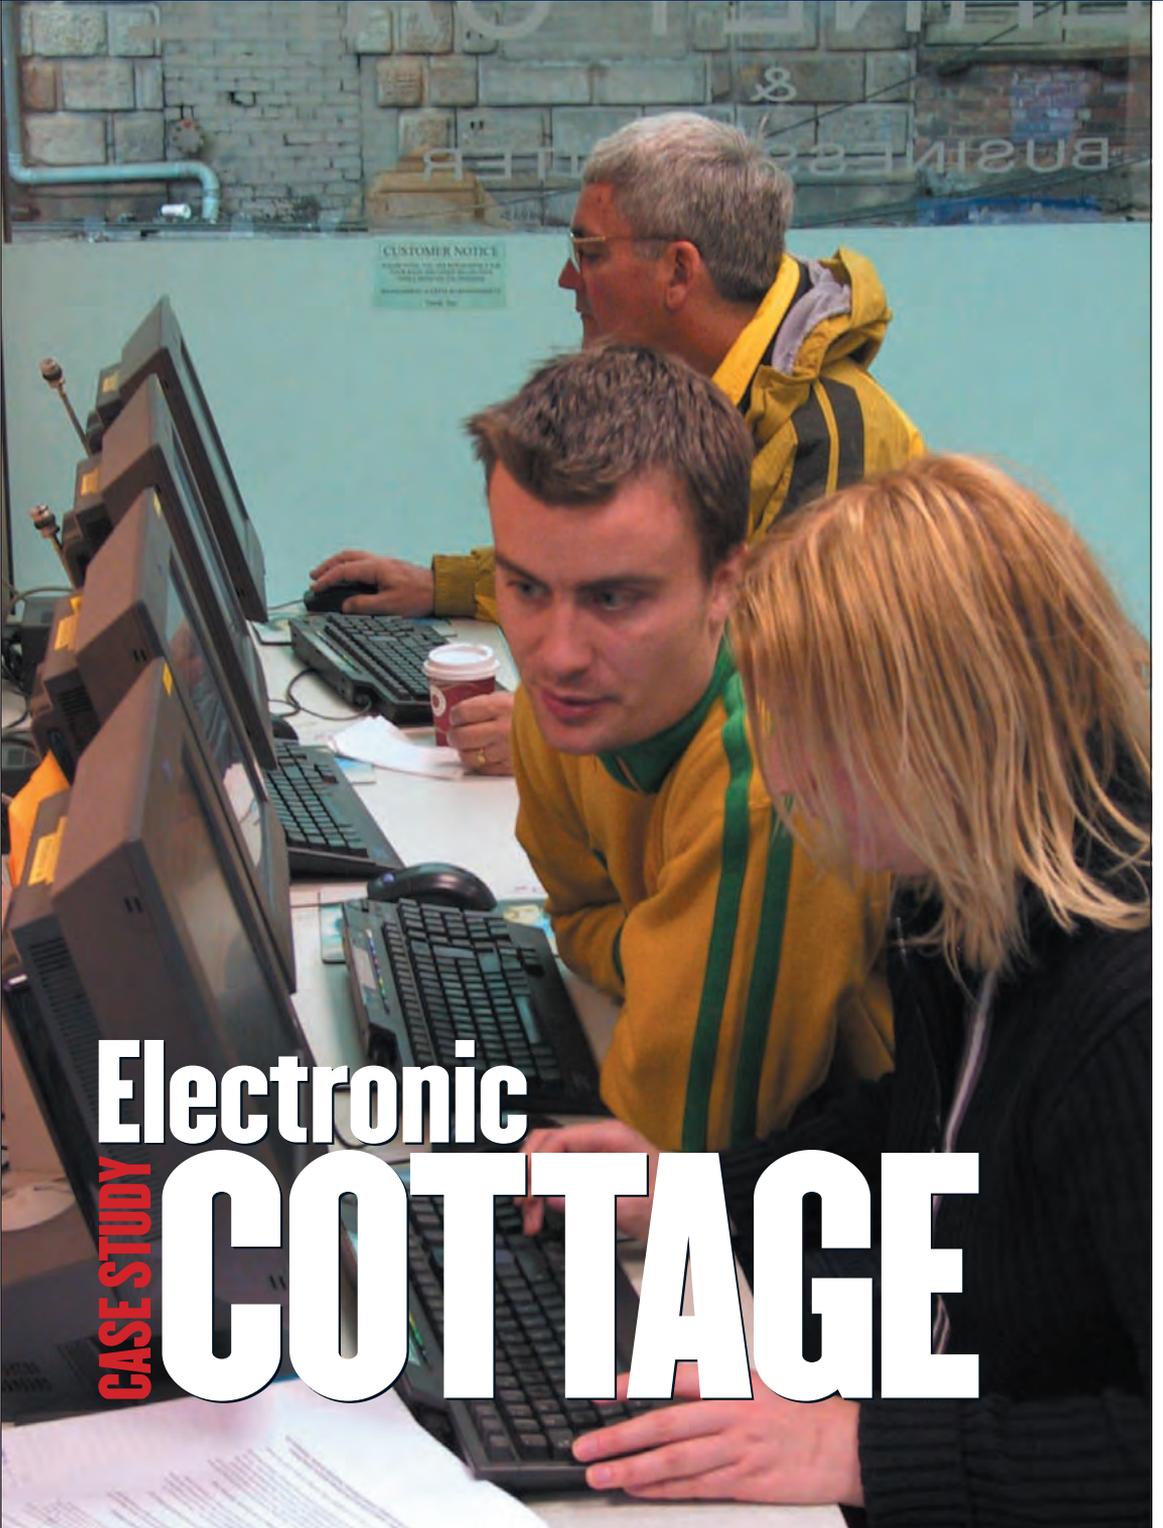
**SURVIVING OPEN SOURCE:**  
 Nick Mailer's thoughts **p12**

**FREEDOM TO INNOVATE:**  
 Patent madness? **p14**

**PORTSENTRY:** Protect your ports from attack **p20**



**CASE STUDY | ELECTRONIC COTTAGE**



**CASE STUDY** **Electronic COTTAGE**

Linux Pro sent **Charlie Stross** to have a latte and a chat with **Subhi Hashwa**, proprietor of **Electronic Cottage**, a small cybercafé in Edinburgh.

**T**he majority of businesses in the UK (and there are several million of them) are small companies. What the American business schools dismiss as “mom’n’pop outfits” form a vital undergrowth that never makes it to the front page of the business section of the newspaper – but we’d be in a real mess if they disappeared. The small business world is where most new ideas come from. One that’s been around for less than a decade is the Internet Café, a public place where you can drink a cup of cappuccino while logging on to your Hotmail account to check your email on the road. It’s a marginal proposition at the best of times; Easy Group’s huge easyEverything cybercafés aren’t reputed to be making an equally huge profit, and while PCs have appeared in the dingier corners of many coffee shops, nobody’s figured out how to go public on a wave of caffeine-craved surfers.

That’s not to say that the humble cybercafé isn’t a viable business model, though. Unsexy it may be, not to mention lacking in opportunities for total world domination and an FT100 listing, but a properly run cybercafé can be a nice little earner and provide a valuable public service. It all depends on what *else* you do – just pouring coffee and providing bandwidth isn’t enough to guarantee success.

Electronic Cottage is distinctive for two reasons: firstly, it’s located in a block of shop units alongside platform 19 of Waverley railway station, the main rail terminus for Scotland’s capital city. Reason number two is that Electronic Cottage’s servers run on Linux, and the client PCs in the café come with a choice of Windows 98 or Linux desktops for the users to choose between. Electronic Cottage hasn’t gone bust; in fact, after just two years they’re looking to open more cafés. What makes the difference?

**Linux Pro:** Is this your first cybercafé?

**Electronic Cottage:** Yeah. We haven’t run one before.

**LXP:** How long have you been open?

**EC:** We’ve been running for over a year and a half – just short on two years, now.

**LXP:** So where did it come from? Why did you decide to open a cybercafé in 1999?

**EC:** Well, we were looking for business ideas, and we took a look at the idea of a cybercafé in Edinburgh because of the tourists and the nature of the location. We aim to provide a quality service and extras that the cafés don’t normally have; we found the location by chance, but we were looking for something like this. I was walking past and saw the empty shop at the end of platform 19, and it looked like a good spot.

*Any retailer will tell you that the three things that make or break a business are location, location, and location. Any number of cybercafés go broke because they open up somewhere out of the way, with no passing trade or in areas where nobody who’s likely to need internet access goes. Platform 19 is one of the two platforms where London-to-Edinburgh express trains stop; if you need to check your email as soon as you’re off the train, this is the place to go and do it.*

**LXP:** Did you look for corporate backing when you set up the café?

**EC:** We’re 100% privately funded. We looked at the possibility of securing sponsorship with an ISP, but the ones we approached weren’t really interested – they’d seen a lot of Internet café’s before and didn’t think much of the field.

*When we visited, the café was surprisingly full; it was a quiet time of day on the run-up to a bank holiday, but still half the computers had customers sitting in front of them. This isn’t automatically the case for a cybercafé.*

**LXP:** How full is the cybercafé, usually?

**EC:** It varies a lot, seasonally, and with time of day; a bank holiday won’t have as many people as a normal day. And there are spikes in demand. On average we reach about 70%.



## CASE STUDY | ELECTRONIC COTTAGE



**LXP:** I notice it looks like you're using IBM client machines – is there any particular IBM connection? And what specification of equipment do you use?

**EC:** For the clients, we settled on machines based on Pentium III/600, which was the best we could get at the time. There are 15 client systems, each with a 15" LCD display, USB, floppy and CD drives – fairly good machines when we opened.

**LXP:** How many servers do you have?

**EC:** We have two servers running Linux, and one running FreeBSD (as a backup system).

**LXP:** Why did you pick Linux for the servers?

**EC:** I had experience running Linux from University. It's reliable and we didn't have the hassle of going through the Windows licensing

**“Publicly embracing Linux as a client-side option has actually brought in new customers.”**

process. We also run Linux in our consultancy business, Electronic Cottage (which the cybercafé grew out of).

*Two Windows NT server licenses, with a fifteen-user license for Internet Information Server, and a licence for Norton Ghost and fifteen client machines would have added considerably to the cost of setting up the café.*

**LXP:** What software do you use on the server systems? How do they work with the clients?

**EC:** We use Red Hat – Red Hat 7.1 and Red Hat 7.2 on different machines. Our main application is written in PHP and runs under Apache; that's the software that controls the café system. The client PCs have a choice of Windows 98 or Red Hat 7.3. All the client computers have static IP addresses on the café network, and the server side software times how long each session runs for, so that users can be billed at the end of their sessions.

**LXP:** What services or software do you provide for your users?

**EC:** Most of our customers are using web browsers or ICQ for chat, as well as web browsing. We also provide printing services, photocopying, CDs, business cards – we're a full business centre. It's a much-needed service at a railway station. I'm told we've had people coming from Livingstone to use our services. Some of our customers, especially European ones, are using *ssh* for remote access to their home or office machines.

*This is a point worth highlighting. Ssh isn't much use to Windows users, but for Linux users on the road, connecting to your home machine via ssh allows you to do things like check the mail securely without going through a sluggish web browser. As Linux becomes more popular as a whole solution – as it already is in Germany – demand for ssh clients will probably rise in cybercafés.*

**LXP:** Are you finding many of the users are using Linux rather than Windows on the clients?

**EC:** After we put out a press release saying people had the choice we got a lot – a lot more people are interested. Desktop users who come from a Windows background read about Linux in a magazine so they come to us to see if they like it. And the set up we have gives them a choice of software.

*This is one of the less foreseeable uses of cybercafés; in addition to network connectivity and coffee, they provide an opportunity for people to hire a computer with new software installed on it, and a risk-free chance to see if it's usable. It's a business opportunity Subhi stumbled across by accident; some proprietors might have actually tried to conceal the Linux-based nature of their network from the customers, in the belief that the unfamiliar would be a deterrent, but in this case being bold appears to have generated additional trade.*

**LXP:** How do you manage your client systems?

**EC:** We re-image them every couple of days. Most of the data storage is NFS mounted from the servers, but we re-image the





operating system partitions to ensure the clients are clean.

**LXP:** What do you use for re-imaging?

**EC:** Instead of Norton Ghost we use partimage – from <http://partimage.sourceforge.net/> – to do the job.

*Partition Image for Linux is a utility that saves disk partitions to an image file. The image file can be compressed to save space and split into multiple files for archiving on removable media. In addition to supporting Linux filesystems such as ReiserFS, ext2/ext3, and UFS, Partition Image directly supports DOS/Windows FAT and Windows NT/2000 NTFS filesystems, and other items (such as Apple's HFS and IBM's JFS). It also has the ability to save and restore master boot records – an important step, because this means you can also restore the boot configuration on machines that boot multiple operating systems. All in all, it's functionally rather similar to Norton Ghost – but the price is a lot nicer.*

*Having taken a look around the nuts and bolts of the café – which seems to be a surprisingly productive use for about 200 square feet of railway platform – we wound up with the obligatory questions ...*

**LXP:** What are your future plans for the café?

**EC:** We're looking for more sites – the right site, of course!

**LXP:** But do you know about BT's recent 802.11 wireless service announcement? Do you think wireless networking is going to hit the cybercafé business model?

**EC:** I don't think it'll have the penetration to affect us, at least a first. We support laptop-carrying people – they can come in and plug in, rather than using our own machines – but 802.11 will take some time to catch on. It doesn't have the percentage of users who use cable to log on; I'd give it three to five years before it's a serious threat. How long will it take for a laptop to ship with it built in?

*Well, let's hope for his sake that Subhi is right. He's not just providing bandwidth, after all – he's providing a full business service in a critically useful location, right next to a major inter-city transport artery. More interestingly, it's a showcase for Linux in business; and publicly embracing Linux as a client-side option has actually brought in new customers.*

*So much for the conventional wisdom that business executives steer clear of software with which they aren't familiar. ■*

**Electronic Cottage makes up for the time you waste when your train is delayed.**





# MEMORY AND uClinux

David McCullough explains the limitations of working without virtual memory, and how to overcome them.

**D**eveloping applications for embedded Linux should be just as easy as developing normal Linux software – that’s part of the appeal of embedded Linux. But while it is true to a large extent, there are certain aspects to consider. Perhaps the most impactful of these is the use of memory, and calls to malloc, and here we will consider such matters with respect to uClinux, as used by SnapGear.

## NOMMU

Firstly, under uClinux we don’t have Virtual Memory (VM.) This means you cannot arbitrarily add memory to an already running

process. As VM is usually implemented using a processing unit called an MMU (memory management unit) you will often hear the term NOMMU when travelling in uClinux circles. Let’s discuss at a very high level how not having an MMU or VM affects malloc.

With VM, all processes run at the same (virtual) address and the VM system takes care of what physical memory is actually mapped to these locations. The virtual memory that the process sees is contiguous but the physical memory it occupies can be scattered about – even on a hard disk in swap!

Without VM, each process must be located at a place in memory where it can be run. In the simplest case, this area of memory must be contiguous, and generally cannot be expanded as there may be other processes above and below it.

Malloc is normally implemented at the low



level using *sbrk/brk* system calls to increase/change the size of a process's address space. The malloc library then manages the extra memory obtained by calling **sbrk()** on behalf of the app. When it runs out of memory it can get more by calling **sbrk()** again, it can also decrease memory using **brk()**, although very few malloc implementations do this. **sbrk()** works by adding more memory to the end of a process (increasing its size). **brk()** can arbitrarily set the end of the process to be closer to the start of the process (reduce the process size) or further away (increase the process size). The standard op of **sbrk()/brk()** is to increase/decrease processes size.

Since *sbrk/brk* cannot increase a processes size under uClinux, malloc needs to undergo some changes at least at the low level in order to work. There are many possibilities here:

- 1 Do not allow processes to allocate dynamic memory.
- 2 Allocate a heap per process that *sbrk/brk* can operate on. The heap size is fixed and is allocated at process startup.
- 3 Let processes allocate memory from the global pool (system wide) of free memory.

Simply put, 1) is just not useful. Too much software assumes the ability to allocate memory and preventing it just increases the effort required to utilise existing apps.

Option 2) has its merits. It limits the amount of memory a process can use (which can be considered a good thing). But it means that the memory for the heap is always allocated, even if it is only used temporarily. The heap space has to be allocated at process start time so that the semantics of *sbrk/brk* still hold and the normal malloc library can be used.

Finally we come to 3). There are pitfalls here. A runaway process can use all of the available memory itself. Allocating from the system pool is not compatible with *sbrk/brk* as they require memory to be added to the end of a process's address space. Thus a normal malloc implementation is no good and a new implementation is needed. The advantage of this method is that only the amount of memory actually required is used. Memory can be returned to the global pool as soon as it's finished with and the implementation can

take advantage of the existing kernel allocator for managing this memory.

Currently method 3 is used by uClinux. The simplest malloc implementation calls *mmap* to obtain memory from the kernel's free memory pool and *munmap* to return it. This gives a very small malloc implementation (1 system call). Problems encountered with such a simple malloc implementation include:

1 The overheads of using *mmap* under uClinux are about 56 bytes per allocation. This turns out to be extremely bad for applications that do a large number of small allocations (like the Zebra routing daemon).

Under uClinux, *mmaps* from user programs are simply allocated from the pool of free memory using *kmalloc*, the kernel allocator, and then chained into a linked list attached to the process. The 56 bytes above comes from the *kmalloc* overhead plus the linked list overheads. Not only is the 56 bytes significant, but a large number of small allocations will result in quite a long list leading to slower deallocations and reallocations.

2 The standard kernel allocator allocates only power-of-two sized quantities up to 1MB, which is inefficient and limiting. To understand the ramifications of this, especially for large

**“Due to power-of-two allocations an application of 130KB in size will actually need 256KB just to run.”**

allocations, consider that a 33K allocation is rounded to the next power of two, 64K!

### Reducing the overhead

Several small malloc implementations have been created to alleviate point 1 by reducing the effect of the 56 byte overhead by allocating larger blocks and then managing those blocks internally for better results.

The kernel allocator has been modified to increase the maximum allocation size substantially. In some cases this is done by a



## FEATURE | MALLOC



kernel config option. This allows for larger allocations and thus larger programs.

An alternative kernel allocator has been added to uClinux that no longer enforces power-of-two allocations, thus eliminating the wastage in allocations quite substantially. This allocator is commonly known as *Kmalloc2* and it can substantially reduce the overheads for memory allocation in a NOMMU environment and increase free memory available.

*Kmalloc2* addresses the allocation limitations by using a power-of-two allocator for allocations up to 1 page in size (a page is 4096 bytes, or 4KB). It then allocates memory rounded up to the nearest page. So for the previous example, an application of 130KB will actually have 132KB allocated to it. A saving of over 124KB on the standard kernel allocator.

*Kmalloc2* also takes steps to avoid fragmenting memory. It allocates all amounts of two pages (8KB) or less from the end of memory down, and all larger amounts from the start of free memory up. This stops transient allocations for network buffers and so on fragmenting memory and preventing large apps from running.

*Kmalloc2* is not perfect, it is quite susceptible to fragmentation of memory, although it could be argued that the standard allocator is more susceptible. *Kmalloc2* works well in practice as the embedded

environments that run uClinux tend to have a relatively static group of long lived applications.

Now we have discussed the kernel level options for memory allocation a little, let's look at the options available to user programs. Due to the issues mentioned thus far, there are quite a few solutions available for use with user programs. Each has its place.

For a start there are a choice of 'libc's' – a topic for another discussion. The choice of 'mallocs' usually depends on the *libc* you are using: *uC-libc* or *uClibc*. Both offer a simple allocator, *malloc-simple*.

*malloc-simple* uses *mmap* and *munmap* to let the kernel actually handle the requests for memory. The implementation is trivial and the code size is negligible, so the cost of including it in an application is very small.

The problem with *malloc-simple* is, as mentioned earlier, that the kernel overhead on an *mmap* based allocation is about 56 bytes. So if you have an app that does a lot of small allocations, its memory usage will be high. *E.g.*, say your program allocates 1000 10 byte amounts – a total of 10000 bytes of memory required. Because of the 56 byte overhead, it will actually allocate 66000 bytes, a 560% increase over what was actually required. *Zebra*, a routing daemon, is an example of an application that suffers quite badly from the small allocation problem as it allocates command data structures as it starts (it has a

## References From the uClinux-distribution sources

Available for download at  
[www.uclinux.org/pub/uClinux/dist](http://www.uclinux.org/pub/uClinux/dist)

### Kernel allocators and mmap implementations:

[linux-2.4.x/mmnommu/slab.c](#)  
[linux-2.4.x/mmnommu/page\\_alloc.c](#)  
[linux-2.4.x/mmnommu/page\\_alloc2.c](#)  
[linux-2.4.x/mmnommu/mmap.c](#)  
[linux-2.0.x/mmnommu/kmalloc.c](#)  
[linux-2.0.x/mmnommu/page\\_alloc.c](#)  
[linux-2.0.x/mmnommu/kmalloc2.c](#)  
[linux-2.0.x/mmnommu/page\\_alloc2.c](#)  
[linux-2.0.x/mmnommu/mmap.c](#)

### uC-libc malloc implementations:

[lib/libc/malloc/alloc.c](#)  
[lib/libc/malloc-simple/alloc.c](#)

### uClibc malloc implementations:

[uClibc/libc/stdlib/malloc](#)  
[uClibc/libc/stdlib/malloc-930716](#)  
[uClibc/libc/stdlib/malloc-simple](#)

Further information on uClinux  
 Further information on SecureEdge™ Development Platforms  
 Further information on SnapGear™ VPN Routers



large number of commands/keywords!)

*uClibc* used to offer a *libsmalloc*, a version of *malloc* that specialises in low overhead small allocations specifically for applications like *Zebra*. This version of *malloc* has since been merged with *malloc-simple* as the extra code required is no longer a significant overhead when using shared libs for uClinux. uClibc offers a few choices as well:

#### **malloc**

A hunk based allocator that appears to have NOMMU support, although it uses *mmap* functionality not available on NOMMU systems. At this point it appears that this allocator will not run on NOMMU systems. It may be relatively simple to fix this, and if so, it should have some potential.

#### **malloc-930716**

This allocator will only work on MMU systems. It relies on the *brk/sbrk* calls. Although these will return a small amount of memory on uClinux systems, it is useless as the sole source of memory for an allocator.

#### **malloc-simple**

The most simple allocator, works on both MMU and NOMMU systems.

Generally the more complex mallocs deliver faster allocations and are more efficient for small allocations. They also add considerable code size to small applications such as one would find in a uClinux environment.

There are a few choices for malloc. Which one should you use? *malloc-simple* is generally a good default malloc to use. Then you can choose a better malloc for the individual applications that need it. As you can see above, your choices are limited if you want something that works out of the box.

Inevitably, despite your choice of kernel/user allocations, you will run out of memory. One of the common problems new users encounter is the "missing memory" problem. The system is showing a large amount of free memory but my application cannot allocate a buffer of size X. The problem here is memory fragmentation, and all of the solutions available at this time suffer from it. Because of the lack of VM in the uClinux environment it is near impossible to fully utilise memory due to fragmentation. This is best explained by

0	100	200	300	400	500	600	700	800	900	1000
-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----										
xxxxx xxxxx	----	xx	-x-	- x--	xx--	- x--	xx--	- x--	xxxx	

example. Suppose your system has 500KB of memory free and you want to allocate 100KB. It is easy to think that this would be possible. However, it is important to remember that you must have a contiguous 100KB of memory in order to satisfy the allocation. Suppose the memory map looks like the table (above). Each character represents approximately 20KB. X marks areas that are allocated or in use by other programs or the kernel.

As you can see, there is 500KB free, but the largest contiguous block is only 80KB. There are many ways to arrive at such a situation. A program that allocates some memory, then frees most of it leaving a small allocation in the middle of a larger free block is often the cause. Transient programs under uClinux can also affect where and how memory is allocated.

The question is often asked, why can't this memory be defragmented. The problem is that we don't have VM and we cannot move memory that is being used by programs. Programs usually have references to addresses within the allocated memory regions and without VM to make the memory always appear to be at the correct address, the program will crash if we move its memory. There is no solution to this problem under uClinux. The developer needs to be aware of the problem and, where possible, try to utilise smaller blocks of memory.

Memory allocation under uClinux, as we have seen, is quite similar to memory allocation under normal Linux but has its own set of quirks and shortfalls. Further progress will no doubt be made on the memory allocation front now that uClinux is enjoying its first shared library implementations. This reduces the requirement to keep the malloc implementation as small as possible as it can live in a shared library. This will lead to larger, but more efficient user level malloc implementations that help to reduce fragmentation and malloc overheads. ■

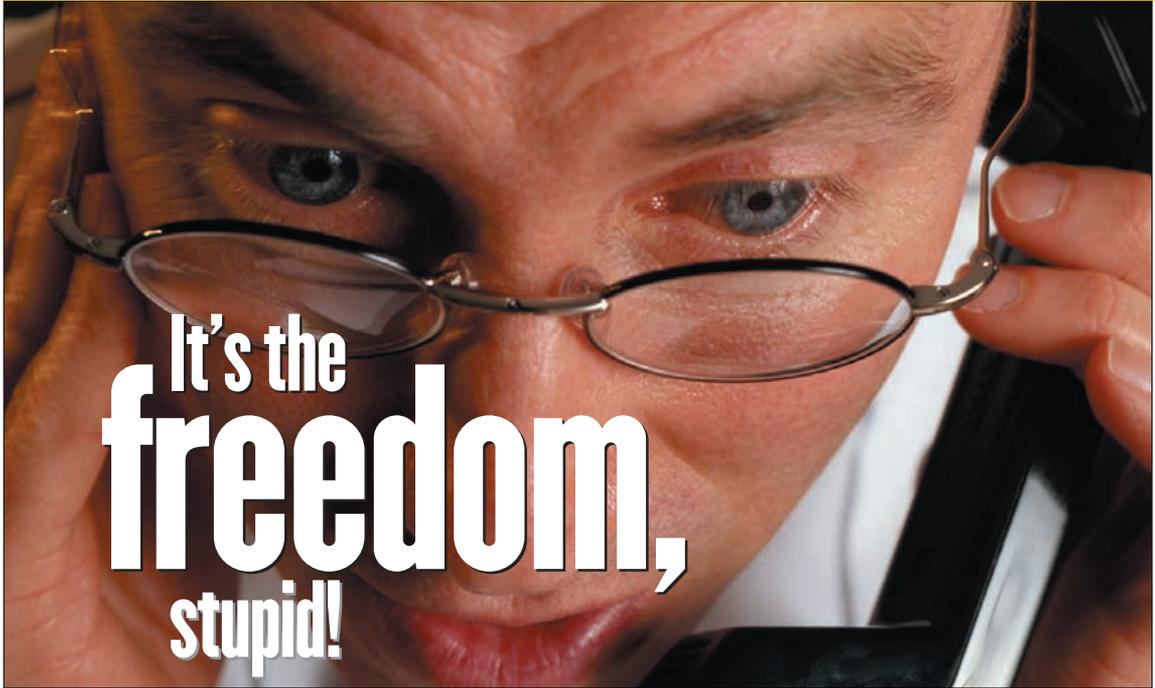


## Thanks

Thanks to Phil Wiltshire for prompting this response. SnapGear and SecureEdge are trademarks of SnapGear Inc. All other trademarks are the property of their respective owners. Copyright 2002 SnapGear Inc.



## FEATURE | OPEN SOURCE BOOM



Nick Mailer of the Positive Internet Company knows what was needed to survive the Open Source boom.



**T**he words went “join us now and share the software, you’ll be free, hackers”. It’s not exactly a business mission statement. In fact, it’s the beginning of the first verse of the whimsical *Free Software Song*, composed by Richard Stallman, the Prophet-like father of the Free Software movement. One might sooner imagine corporate boardroom walls sporting posters of Che Guevara than basing their IT strategy on software emerging from this movement. “When we have enough Free Software at our call, hackers, . . . we’ll kick out those dirty licenses, ever more..” the ditty continues. Techno flower-power at the entrance of the big-tented, globally spun third millennium? Surely not. “Share the software?” Charming, perhaps; naïve, certainly, but no more viable a world view than proclaiming “ban the bomb” or, indeed, teaching the world to sing in perfect harmony.

Free Software’s being the basis of an active

revolution in business ethics and economics comes, therefore, as something of a surprise.

For companies in the know, particularly those with an interest in informatics, Free Software simply makes sense. Free Software, mind you, rather than its business-friendly, diluted incarnation, “Open Source.” Eric Raymond created the latter phrase with Bruce Perens, both software developers and social commentators. It was meant to be a business-friendly gloss upon (or castration of, as some wags put it) Free Software. It was designed to coax the weary and timorous corporate executive out of the proprietary-software bolt hole with promises of efficient methodologies, quick bug-fixing and, no doubt, the extraction of algorithmic light-beams from object-oriented cucumbers.

To an extent, the marketing spin worked. Some organisations, like Netscape and IBM, began to realise that, yes, the methodologies described in such works as Raymond’s seminal



**OPEN SOURCE BOOM | FEATURE**

*The Cathedral and the Bazaar*, did have business merit. The money was rolling in, the dotcoms were booming and, heck, if dress-down Friday, roller-blading through the corridors and this new-age “Open Source” thing were going to keep the employees happy and the venture capitalists continuing to capitalise ventures, then so be it. Like the word “democracy”, though, the term “Open Source” found itself easily abused. In the feeding frenzy that followed its minting, some were less picknickety than others about visiting the Open Source Institute’s website to find out its actual definition. It seemed almost as if some software companies put their same proprietary software in a newly pastel-shaded box and decide that it was, therefore, “Open Source”. Others would let you see bits of their source but not change it. Others still would allow you to change bits of their source, but not release them. Everyone wanted to join the “Open Source” dance, but very few knew the moves or cared about the true state of the ballroom. As long as the music was playing, who cared?

But then, of course, the music stopped playing. The Venture Capitalists sobered up: loans were foreclosed; investment was withdrawn; whole industries were devalued. It all began to feel like a barely-remembered night of drunken excess. Unfortunately, the term “Open Source” had been bandied about by some of the heady companies that were now just embarrassing footnotes to this corporate booze-up. So, as with so many of the brave new industries keen to be seen to adopt it, “Open Source” became tarnished. Even Bruce Perens disowned the phrase and went back to talking about Free Software.

**The survivors**

But there were survivors of the dotcom massacre. Those businesses that had rode through the boom without going bust had usually borrowed less money, produced less hype and had been less credulous. In relation to Open Source software, many realised its true strengths were the characteristics so downplayed by the Open Source Institute. It wasn’t the methodologies of the software’s production that suddenly appeared interesting

Rather than the engineering efficacy it was the lustre of the liberty that now shone brightly.

The companies that had survived the bust were those whose founders valued freedom: from grasping venture capitalists and from someone else’s business plan. The survivors had known the value of being unencumbered. A truly unencumbered company worries whether its software comes with strings attached. An organisation in control of its own destiny cares whether a monopolist is about to play cat-and-mouse games with licensing structures. Ever more organisations are beginning to realise that “free as in beer” isn’t the issue. Most of these companies were founded with a libertarian, entrepreneurial

**“Companies that survived the bust were those that valued freedom”**

spirit. The same spirit is revealing that the software, the very conduit of the business, can only be future-proof if it is “free as in speech”. And so, with a certain tincture of Boston Tea Party glee, businesses all over the world are, indeed, kicking out the “dirty licenses”. In the end, it isn’t the Cathedrals. It isn’t the Bazaars. It’s the Freedom – stupid! ■

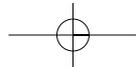
**A company that recognises the value of freedom for businesses.**

The screenshot shows the website for The Positive Internet Company Ltd. At the top, there's a navigation menu with links: about us, managed servers, dolphin server, contact us, our clients, shared hosting, domain names, support, latest news, order, fun, and jobs/positive. Below the menu, there are three main service sections: Fully Managed Servers, Shared Hosting, and Dolphin Server. Each section lists various features and benefits. At the bottom, there's contact information including a freephone number (0800 316 1006), the company name and address (The Positive Internet Company Ltd, Positive House, 24 Broadway, London W1A 2SU), and an email address (good@positive-internet.com).

FEATURE | PATENTS

# FREEDOM to innovate

Software patents aren't just a threat to Free Software, they're a danger to all software development, as Richard Smedley discovers.



**T**he European Community seems set on course, after intensive lobbying, to introduce software patents. These are not patents on particular programs – programs are already well protected by copyright law – but on the *ideas* which they contain. Patents have much in common with a lottery – they are sold by reference to the big prizes on offer. However, as with a lottery, most of us are just small losers.

Richard Stallman (popularly known as RMS) has been campaigning on copyright issues for two decades, as part of his founding of the Free Software Foundation (FSF) and initiating the GNU Operating System. Lately he has also brought to Europe his campaign against software patents as they threaten the development and dissemination of Free Software. However the case he makes is that software patents threaten *all* software development.

When RMS was invited to speak at Aberystwyth by the Department of Computer Science at the University of Wales, during a recent visit to the UK, *Linux Pro* jumped at the chance to hear the talk and catch up with RMS afterwards for a further explanation of the issues.

RMS makes a clear distinction between patents and copyright. “Patents have nothing to do with copyright. They are two different areas of law which evolved separately with no relation to each other, that apply to different things, and every detail’s different. Never lump them together

“There’s a term I’m sure you’ve heard – ‘intellectual property’ – which is a big, catch-all generalisation which lumps together copyright, patents, trademarks and various other less commonly encountered things. They’re all different from each other. If you try to think about them all together you won’t understand any of them. All your conclusions will be simplistic. You’ll be wrong in any factual conclusion that you try and reach. And when you try and look at any public policy questions you won’t be able to because those arise from the details – that are different. When you look at this broad abstraction, ‘intellectual property’, you can’t even see any of the public policy

issues raised by copyright or raised by patents, or the issues raised by trademarks, because they’re all different issues. If you do try and think about public policy at that level, all you can do is choose among a few simplistic over-generalisations – none of which is any good.

“So the first step for thinking about copyrights and thinking clearly about patents is to disconnect them – make sure that they’re completely separate in your mind – and never use the term ‘intellectual property!’”

## Winners and Losers

Lottery tickets – most bring no benefit, a few bring lots of money. The winnings in this case are promoted by patent lawyers, but the real victims are developers: Want to write a

**“The holder of the patent wanted 5% – what if you need to license twenty patents for your app?”**

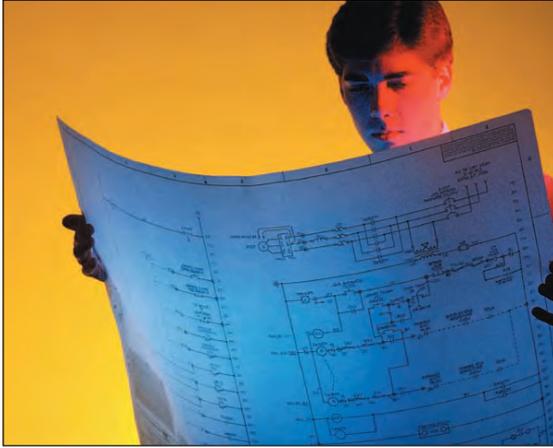
program to get something done? You’d better start by looking at what patents are held in the field. One problem is that patents pending are often kept secret until after they are granted – this may be eighteen months down the road when you have already written and published your code. Now you are in trouble.

RMS gives the example of *Compress*, which, in 1984, was “distributed for various Unix systems. At the time there was no problem,

## The differences Patents bear no relation to copyright

Copyright covers details, not ideas. A patent is a restriction on using a certain idea, copyright is a restriction on *copying*. A patent is an absolute monopoly on using a certain idea – even if you also thought of the same idea yourself. Copyright lasts up to 150 years. Patents are currently limited to 20 years. Copyright happens automatically but patents must be applied for, and are expensive. Copyright has no threshold (for interest or originality) and covers everything published. Patents must be original, not an obvious idea (at least in theory, though they often aren’t).





but in 1985 a patent was issued in the US on the LZW data compression algorithm used in the *Compress* program. And the patent holder didn't do anything about this immediately – the patent holder (Unisys) figured 'let's get everybody habituated to using this, then we'll stick them up and they won't be able to stop'. So several years later they started to demand royalties for the use of *Compress*, and we in the GNU project started to look for some other way to do data compression.

"I found somebody who'd written a program already and decided to contribute it to us. A week before it was going to be released I happened to see in the *New York Times*' weekly patent column, which said somebody had got a patent for inventing a new method of data compression – it turned out to cover the program we were going to release, so that program was killed, never got anywhere. We

eventually found someone else who came up with *another* algorithm and we released the program generally known now as *Gzip*."

Of course you may not find the patent, even if it has already been issued. Relational searches will turn up many patents in the field of your proposed program, but the language used in the patents means that it is not always clear what they are covering.

In the USA, a patent on a technique for compiling formulas to object code, which was applied for in 1970 but not granted until 1983, was used to try and squeeze money from every company making a spread sheet. The idea covered was topological sorting using the natural order recalculation algorithm – which recalculates each cell *after* the cells upon which it depends are changed, so that the sheet is always up to date. The patent made no mention of topological sorting or natural order recalculation, and, even if you know what you are looking at, it still looks like the wrong patent.

An Australian government report (which concluded that software patents were a bad thing) discovered that few engineers read patents – "I can't recognise my own ideas in patentese," said one – negating the disclosure benefit originally cited for patents.

## Dealing with patents

RMS outlines three approaches for dealing with software patents, all of which have disadvantages.

### 1 Avoidance

Simplest is to avoid the patented feature altogether. In fact users of the *xyWrite* word processor once actually received a

## Hypercard "What did I patent?"

**RMS:** "In the 1980s an engineer got a couple of patents which are very difficult to describe in a nutshell. However when he saw *Hypercard*, he didn't think it was what he had patented. Then his lawyer told him that there was a way of interpreting

the patent such that *Hypercard* infringed it, so he threatened to sue Apple – who ignored him. He then started threatening to sue Apple customers. At this point Apple settled. "So you're going to have to spend a lot of money finding out what these

patents do to you, and ultimately you'll get an answer like this: If you do something in this area you're almost certain to lose. If you do something in this area there's a substantial chance that you'll lose. If you really want to be safe then stay out of this area."



## Lottery winners ...and it's not the starving inventor in the attic

In fact the only winners of this lottery are companies large enough to buy themselves a million tickets. IBM may make some money from royalties on their patents, but they make far more from getting access to the patented technology of others through cross licensing. In 1990, in an article for *Think* magazine (<http://lpf.ai.mit.edu/Links/prep.ai.mit.edu/ibm.think.article>), IBM claimed that using their (then) 9000 active US patents gave them of

the order of *ten* times more profit from the benefits of cross licensing than from royalties.

"This is a game that the rest of us can't play. Few small businesses get even a dozen patents," says RMS. It exposes the myth of the "starving genius in his attic. Some big company will take all his business away. Development in a high-tech field doesn't happen in attics."

Of course if an individual did come up with a great idea, and patent it,

she would probably find that someone had patents on other ideas in her program, and she would have to cross license with them, losing the protective monopoly that patent proponents claim that they give.

It is these big companies who benefit from cross licensing, and who now have large amounts of value stored in patents, who have been lobbying hard for software patents, not the smaller companies or individuals you may expect.

'downgrade' for the product. The developers were threatened over a feature which allowed user-defined abbreviations for long words or phrases. After unsuccessful negotiations with the patent holder they decided to simply remove the feature. Modern word processors have hundreds of features, any one of which could be patented alone, or in certain combinations with other features – how many features can you leave out of your program?

There was certainly no way around the public key encryption patent, which largely blocked public key encryption in the USA for ten years. Sometimes market inertia is just too great to shift. GIF files, compressed with the LZW algorithm, were the standard ten years ago when the FSF urged people to switch to the patent-free PNG format. As not all browsers supported PNG, webmasters didn't switch, even after Unisys (the LZW patent holder) started threatening browser companies and people with websites.

### 2 Licensing

The patent holder does not have to offer you a licence, but if they choose to do so you may not be able to afford it.

The holder of the natural order recalculation patent wanted 5%. What if you need 20 licences – or 21? Economists say that two or three licences at these (not untypical) rates would make any project economically unfeasible. Needless to say that licences

working on even the tiniest percentage are unfeasible for Free Software as the controls needed to ensure individual compliance would render the software non-free.

However it is in the field of patent licensing that we see the real winners of the patent lottery (see the *Lottery winners* box). Exposure of this cross-licensing myth also shows that patents do not protect the small inventor from large corporations, as has often been claimed by patent advocates.

**“Imagine 18<sup>th</sup> century patents on music – Beethoven would have been unable to write a symphony”**

### 3 Litigation

Apologists for patents often say that however many bad patents are granted, they can always be brought before a judge to be settled reasonably in a court of law. Qualcomm were recently in court and had to pay out \$13 million – most of which went to the lawyers on both sides. European courts are a little less expensive than their American counterparts – but how many companies can run the risk of losing even one patent battle a year?





## Tread carefully

If you face several patents covering your project you may try the above three options for each patent, but sometimes you might get stuck. Getting past each one is like crossing a minefield: "To develop a piece of software you may not tread on a patent with each step, but to cross the whole field is not so easy!"

So why is software so different when it comes to patents? Other fields of endeavour seem to cope with the system.

## Complexity

The main factor here is the difference between the physical world – with all its various unpredictable constraints – and mathematics, where things are what we define them to be. As RMS says, "When I put an if statement inside a while statement I don't have to worry

about whether the loop will run at the wrong frequency and the if statement will start vibrating and will fracture. I don't have to worry about whether with this vibration it will rub away part of the while statement!" The sheer predictability of code – the way it stays the same, and is so easy to manufacture and reproduce – means that programs often contain hundreds of thousands of 'components'. Designing a physical system with a million components is a major undertaking.

At the opposite end of the spectrum a biotechnology product may be protected by a single patent on the molecule. Another product would be covered by a different patent. Software is also much cheaper to design – a program cannot cover the cost of several patents in the way that a similarly complex physical system may be able to.

Patents are an artificial system to (allegedly) 'incentivise' ideas. However ideas are not in short supply when it comes to creating software – they tend to flow as a byproduct of getting developers together to do the hard part of the job – writing the code.

## Making allies

As we said earlier, software patents affect all coders – not just developers of Free Software, so while Free Software organisations have been amongst the first to recognise the threat, there is a broad coalition building against them. Indeed large parts of the French and German governments have come out against the UK Patent Office-led attempt to introduce software patents here.

The League for Programming Freedom (<http://lpf.ai.mit.edu/>) has long campaigned against software patents. Recent developments have seen the formation of European groups to counter the threatened new laws. The Foundation for a Free Information Infrastructure (<http://ffii.org/index.en.html>) wishes to "protect innovation and competition in software industry". Free Patents ([www.freepatents.org](http://www.freepatents.org)) also has useful information on its website, and invites you to join its mailing list. <http://liberte.aful.org/mailman/listinfo/patents>

While Linus Torvalds may have come out

**“Red Hat claim they won't enforce the patents where they are used in copyleft Free Software”**

## Further information

Gareth Bowker organised RMS's talk, which was sponsored by the Department of Computer Science at the University of Wales, Aberystwyth: [www.aber.ac.uk/compsi](http://www.aber.ac.uk/compsi)

A recording can be found split into three convenient chunks at: [www.6809.org.uk/audio](http://www.6809.org.uk/audio)

If you dump them to wavs and speed them up by about 2% you can even get the largest one onto an 80m audio CD (thus making a nice 3 CD set gift for someone ;-) There's a link there to the whole

thing as one big (90MB) file residing on [ftp.bbc.co.uk](ftp://ftp.bbc.co.uk).

The UK's Association For Free Software ([www.affs.org.uk](http://www.affs.org.uk)) are part of the campaign against software patents.

To make yourself heard, sign the petition: <http://petition.eurolinux.org/>

Lastly, to show that no battle covers new ground, visit: <http://odur.let.rug.nl/%7Eusa/P/tj3/writings/brf/jefl220.htm> for Thomas Jefferson's letter, *No Patents on Ideas*.





strongly against software patents, not all distributors of GNU/Linux have been so clear on the matter.

### Patented kernel?

SELinux (<http://www.nsa.gov/selinux/>) is a distro produced by the U.S. National Security Agency, based on the Linux Security Module architecture. With no 'root' user, it provides mandatory access control features to protect parts of the system from each other. Any compromise is highly limited in the damage it can do to the rest of the system. SELinux is under the GPL. However the implementation of SELinux was done by Secure Computing Corporation (SCC), who used a technology that it calls type enforcement. SCC has a patent on this technology.

Recently on Linux Weekly News (<http://lwn.net/Articles/2379/>) an SCC employee was quoted as saying "SELinux includes Type Enforcement technology developed and patented by the Secure Computing Corporation, who still holds rights to all commercial use of the technology. Before a colo company, or anyone else uses the technology commercially, it will be necessary to negotiate a license with Secure Computing. If anyone wants to do so, I can help get the ball rolling with our Legal and BD folks." SCC have said that they plan to make their position clear, so watch the news sites.

Far more controversial has been the involvement in patents of noted GPL fans Red Hat. Whilst consistently condemning patents in its published policy, it has emerged that Red Hat themselves have taken out two software patents (filed last August). They argue that they are prudently building up a portfolio for defensive purposes – giving themselves the ability to cross license, and thus avoiding being put out of business by someone who holds a patent on an algorithm used in Free Software.

Red Hat claim that they will not enforce the patents where they are used in copylefted Free Software. See [www.redhat.com/legal/patent\\_policy.html](http://www.redhat.com/legal/patent_policy.html) for the full explanation.

### Conclusion

In a system as complicated as software development there is now mathematical modelling to show that when you have a field with incremental innovation, having a patent system can obstruct progress: [www.researchoninnovation.org/patents.pdf](http://www.researchoninnovation.org/patents.pdf)

Patents affect us all – even users of 'infringing' programs can be sued (see *Hypercard* box). They make extra work for developers, with no benefit to the small or medium sized business. "Software patents are a new kind of bureaucracy that can tie up every computer user," says RMS, and the last thing that European businesses need is yet more bureaucracy. ■

## WHAT DO YOU THINK?

Do you have strong views for or against software patents? We're interested in your opinions. Send us your feed back to [linuxforum@futurenet.co.uk](mailto:linuxforum@futurenet.co.uk), and put "Patents" in the subject line.



# Installing and running PORTSENTRY

In this month's security tutorial **Charlie Stross** explains how to install PortSentry and keep the bad guys out while you're on the Internet.



**P**ortSentry is a program that actively responds to attempts to hack into a computer via the Internet. It's free (as in beer), it's reasonably easy to install, and it can detect attempted attacks and lock the attacker out, while notifying you of where they come from.

*PortSentry* comes as standard with most newer Linux distributions, but you can obtain the source code from Psionic Technologies, [www.psionic.com](http://www.psionic.com). While the licence terms do not meet the Open Source Definition, the source code is provided and you're allowed to use the product for your own internal purposes.

Two versions of *PortSentry* are available: version 1.1 and a beta of version 2.0. The major changes in version 2.0 are that it detects stealth scans (port scans using malformed packets that are designed not to result in a fully-opened socket connection, and thereby avoid normal logging mechanisms). Version 2.0 also eliminates the "classic" port-binding mode. In this tutorial we're going to cover installation of *PortSentry 1.1*, because 2.0 is not yet a stable product.

## Installing PortSentry

To start with, you need to register to download the source from [www.psionic.com](http://www.psionic.com). The total source code package is approximately 45KB in size, and once you've downloaded it you will need to unpack it – as root:

```
tar xvf portsentry-1.1.tar.gz
```

### cd portsentry-1.1/

Compilation is extremely easy, if you're willing to accept the default settings. **make linux** will build the *PortSentry* system on most Linux systems, or **make debian-linux** if you're running Debian or a debian-based system (such as Corel's Linux distribution or Stormix).

You may want to customise some variables in the `portsentry_config.h` file. These include `CONFIG_FILE` (the path to the *PortSentry* configuration file `portsentry.conf`), `WRAPPER_HOSTS_DENY` (the path name to your TCP wrapper `hosts.deny` file, usually `/etc/hosts.deny`), `SYSLOG_FACILITY` (the syslog facility to use), and `SYSLOG_LEVEL` (what syslog alert level to tag messages with). However, the default values are fine for most Linux systems.

When you've built *PortSentry*, you need to do the basic installation:

```
$ make install
```

```
Creating psionic directory /usr/local/psionic
```

```
Setting directory permissions
```

```
Creating portsentry directory
```

```
/usr/local/psionic/portsentry
```

```
Setting directory permissions
```

```
chmod 700 /usr/local/psionic/portsentry
```

```
Copying files
```

```
cp ./portsentry.conf
```

```
/usr/local/psionic/portsentry
```

```
cp ./portsentry.ignore
```

```
/usr/local/psionic/portsentry
```

```
cp ./portsentry /usr/local/psionic/portsentry
```

```
Setting permissions
```



## Portscanners explained Rattling the door handles

Port scanning is a technique used by crackers to gain access to servers. Just about every Linux machine qualifies as a server – it's how most Linux software works. Some versions of some server programs have known bugs; in some cases, by feeding a buggy server a carefully contrived lump of bad data, a cracker can cause it to execute programs of their own choosing on the server machine. Many server programs need to run as root, and if a cracker can run a program as root they can eventually gain unlimited access to the machine (by, for example, giving themselves a user ID and password with root privileges).

Port scanning is often the first stage of an assault on your computer. Every box on the Internet is characterised by its own IP address – a series of four numbers that uniquely identify it as a destination. (The human-readable domain names we use are translated to the underlying IP addresses by the Domain Name System, a distributed database that exists to make the Net easier for humans to use.) Because computers may be using several

services at once, in addition to specifying an IP address as the destination, each network packet also needs a port number: ports are simply a number (from 0 to 65535) corresponding to a particular service session. In addition, connections require a connection type – either TCP (serialised) or UDP (broadcast) connections are available, and a port 80 TCP connection is not the same as a port 80 UDP connection.

Many network services are provided by servers that listen for incoming requests on “well-known” ports; for example, all SMTP (email) connections run on TCP port 25, and almost all HTTP (web) servers listen on TCP port 80. Almost all the main services have port numbers below 1024 – these are “privileged” ports and, on Linux, non-root programs can't use them.

Port scanners are programs that, for a given machine with a set IP address, scan each port – usually sequentially, probing port 49, then port 50, then port 51, and so on. If a server is running on the target, the portscanner will connect to it and try

to identify the server program that's running. It's the equivalent of walking around a building testing the doors and windows to see if they're unlocked. Innocent users, in contrast, usually connect to a few well-known ports and don't test ports sequentially.

Some portscanners are clever enough to run “stealth” scans. Rather than going through all the steps involved in opening a network connection, they proceed halfway, then stop. The server times out the connection after a few seconds, usually without logging the connection attempt or handing the connection over to a server program. Other “smart” scanners probe ports at random, or probe whole IP address ranges in parallel (rather than hammering on a single computer).

Other portscanners pick port numbers at random over a period of time. *PortSentry* remembers hosts that connected in the past, and activates once a user-defined threshold for suspicious activity has been crossed even if ports are not being scanned sequentially.

```
chmod 600 /usr/local/psionic/port Sentry/
port Sentry.ignore
```

```
chmod 600 /usr/local/psionic/port Sentry/
port Sentry.conf
```

```
chmod 700 /usr/local/psionic/port Sentry/
port Sentry
```

Edit `/usr/local/psionic/port Sentry/port Sentry.conf` and change your settings if you haven't already (route, etc). WARNING: This version and above now use a new directory structure for storing the program and config files (`/usr/local/psionic/port Sentry`). Please make sure you delete the old files when the testing of this install is complete.

*PortSentry* installs everything under `/usr/local/psionic/port Sentry`. Before you can run *PortSentry*, you need to edit the

files `port Sentry.conf` and `port Sentry.ignore` in this directory.

### Configuring PortSentry

You now need to edit two files: `port Sentry.ignore` (a list of hosts and networks not to block) and `port Sentry.conf` (the master configuration file). These are text files that can be edited using *vi*, *emacs*, *pico*, or any other text editor.

If you have a single computer connected to the Internet via dialup or cable, you can probably use the `port Sentry.ignore` file without modifying it. If you are installing *PortSentry* on a firewall or gateway machine, things are different. This file is intended to list hosts and networks that you never want to block –





typically, any local networks behind your firewall belong here. Suppose you are using the private-use-only 192.168.1 class C network, with a netmask of 255.255.255.0, for your internal LAN. You need to add a line to `portsentryignore` like this:

```
192.168.1.0/24
```

The format is IP address/netmask, with the netmask usually specified as the number of bits to set. The effect of the line above is to tell *PortSentry* not to block any machines on this network – so use it with extreme caution!

Note: `portsentryignore` comes with the loopback address (127.0.0.1) set to ignore. Don't change this, or you'll probably be sorry.

Next, we need to edit `portsentry.conf`. This file is heavily commented (comments are lines beginning with a hash symbol). It specifies roughly in order: which ports *PortSentry* should monitor, what stealth scan detection strategies to use, where *PortSentry* other configuration files are located, and how to respond to an attack. Let's look at each of these in turn:

### Port configurations

The `TCP_PORTS` setting specifies a list of ports to monitor. Three pre-canned settings are provided – “really anal”, “aware”, and “bare bones”. Note that “really anal” will prevent you from running X-Windows on the *PortSentry* machine, and none of these settings affect stealth scan detection.

**“PortSentry recalls hosts that connected in the past, activating even if not scanned sequentially.”**

When running in normal mode (using the settings in “Port Configuration”), *PortSentry* binds to each specified port and listens for an incoming connection. When running in advanced modes (of which there are a few), *PortSentry* doesn't bind, but monitors a range of ports at the socket level. The key variables

to tweak here is “ADVANCED\_EXCLUDE\_TCP” and “ADVANCED\_EXCLUDE\_UDP”. These should contain comma-separated lists of those ports below the value of `ADVANCED_PORTS_TCP` or `ADVANCED_PORTS_UDP` (usually 1024 in both cases) which you are running servers on. For example, if you are running a web server on TCP port 80 you really want to make sure that `ADVANCED_EXCLUDE_TCP` contains “80”, otherwise HTTP requests for port 80 will be picked up by *PortSentry* rather than being passed through to your web server.

(You switch on the stealth scan detection modes by starting *PortSentry* with the `-stcp` or `-sudp` flags, and advanced scan detection using the `-atcp` or `-audp` flags. Normally, however, you don't need to use these modes unless you specifically expect to come under heavy attack by crackers with stealth scanners.)

### Response Options

From here on, the `portsentry.conf` specifies how we go about responding to probes.

Firstly, we set `BLOCK_UDP` or `BLOCK_TCP` to 1 if we want to block scans; if all we want to do is log them, we set these variables to 0. Next, we have a choice of responses. We can chose to route packets from an attacker to a non-existent host by manipulating our kernel routing table, add the attacker to the `/etc/hosts.deny` file (so that *inetd* will not start servers in response to requests from the attacker), run an arbitrary external command, or use *ipfwadm*, *ipchains*, or *iptables* (depending on your kernel version – 2.0, 2.2, or 2.4) to drop all packets from the offender.

Probably your best choice is to use packet filtering; a number of examples are given in the `portsentry.conf` file, and you can simply uncomment the one appropriate to your packet filtering mechanism (which depends on your kernel). For example, on a 2.4 kernel, uncomment the line:

```
KILL_ROUTE="/usr/bin/iptables -I INPUT -s $TARGETS -j DROP"
```

If your kernel doesn't support packet filtering you may want to add a routing table entry that rejects packets from an attacker:



```
KILL_ROUTE="/sbin/route add -host
$TARGET$ reject"
```

Or for older kernels define a bogus, non-existent gateway so that the packets bounce off in search of a non-existent host:

```
KILL_ROUTE="/sbin/route add -host
$TARGET$ gw 333.444.555.666"
```

Adding the offender to your `hosts.deny` file is good practice; it stops them from connecting to servers launched by `inetd`. Two example lines are provided with syntax for the `hosts.deny` file on old style and new style machines. For example:

```
KILL_HOSTS_DENY="ALL: $TARGET$ :
DENY"
```

(This means that all services are denied to the attacker, `$TARGET$`.)

## Scan trigger

Finally, you want to set the `SCAN_TRIGGER` variable (near the end of `portsentry.conf`) to 2, to reduce the risk of false alarms. This is the number of connections from another host that is required before `PortSentry` assumes that it is under attack – the default setting of "0" puts `PortSentry` on a hair-trigger, and a setting of 2 makes allowances for an innocent mistake (but not a scan).

## Running PortSentry

To start `PortSentry` in standard mode, run it once with the `-tcp` flag and once with the `-udp` flag, thus:

```
/usr/local/psionic/portsentry/portsentry -tcp
/usr/local/psionic/portsentry/portsentry -udp
```

(You can add this to a startup script in `/etc/init.d/rc.d` if you feel confident about editing init shell scripts.)

You should see something like this in `/var/log/messages`:

```
May 3 14:11:04 icebook portsentry[2298]:
admiral: Psionic PortSentry 1.1 is
starting.
May 3 14:11:04 icebook portsentry[2299]:
admiral: Going into listen mode on TCP
port: 1
May 3 14:11:04 icebook portsentry[2299]:
admiral: Going into listen mode on TCP
port: 11
```

## How to lock out the bad guys

One way of reducing vulnerability is simply not to run servers that you don't need. Another is to use a tool like `PortSentry`, which grabs any unused ports and waits behind them to trap intruders. Note, however, that you shouldn't block all the unused ports on your PC. Every outgoing connection you make – e.g. looking at a web page or sending email – requires the temporary use of a port number on your own box, so telling `PortSentry` to monitor all unused ports will prevent you from using your own system as a client.

`PortSentry` can respond to attacks in several ways. Firstly, if all you want is a passive warning that you're under attack, `PortSentry` can record the incident in your system logfile via `syslogd`. `PortSentry` can also attempt to take defensive action, in three ways. It can add the attacking host to your `/etc/hosts.deny` file, so that any services started by `inetd` (the internet daemon) using the `tcpd` wrapper will deny

connection attempts from the attacker. Alternatively, it can update your kernel routing tables, so that packets from the attacker are forwarded to a non-existent host (and therefore time out). This is not ideal – a persistent attacker generates many open connections to clog up your system – but it works on most Linux systems. The third strategy only works on machines that have a packet filtering system such as `ipchains` configured: `PortSentry` can configure the packet filter to drop all packets from the offending host.

One important point to note is that if you're running `PortSentry` on a gateway machine – e.g. a firewall hooked up to a cable modem, providing masquerading services for your home LAN – you must tell it to ignore scans originating from localhost and from addresses on your LAN. Or you run the risk of accidentally locking out your own PC if, for example, you try to run `Netsaint` to look for local vulnerabilities.

[and so on]. When someone scans you, you'll see something like this:

```
May 7 13:13:43 icebook
portsentry[26991]: attackalert: Connect
from host: raq981.uk2net.com/
213.239.59.58 to TCP port: 4001
May 7 13:13:43 icebook portsentry
[26991]: attackalert: Host 213.239.59.58
has been blocked via wrappers with
string: "ALL: 213.239.59.58"
May 7 13:13:43 icebook portsentry
[26991]: attackalert: Host 213.239.59.58
has been blocked via dropped route using
command: "/sbin/route add -host
213.239.59.58 reject"
```

That's all! ■

