

FREE 24-PAGE LINUXPRO SUPPLEMENT TURN OVER NOW!

LINUX

www.linuxformat.co.uk**FORMAT**

YOU, COPYRIGHT AND THE LAW

Know your rights! Keep it legal with our 7-page copyright exposé **p46**

**THE GIMP
SECRETS
REVEALED**

Create great graphics – we show you how **p72**

**JUDGE FOR
YOURSELF!**

Seven top sound apps go head-to-head in our roundup – and they're all on your coverdiscs **p40**

NEW SERIES!**Learning Python**

Start scripting now – tutorial series starts this issue! **p80**

X Apps on Mac OS?

The Fink project – porting your favourite software to Mac OS X **p54**

SVG versus Flash

The Scalable Vector Graphics image format explained **p62**

Office without Microsoft
SuSE's new Office Desktop edition targets Windows users **p20**

DVD issue also available Printed in the UK

LXF38 MARCH 2003 £5.99

 **the future network**

Media with passion 9 771470 423019 03

(£5.99 outside UK and Republic of Ireland)

THE UK'S BEST-SELLING LINUX MAGAZINE

Rights and wrongs

Cult cinema may not be your bag, but indulge me a moment. There's a film called *The Return of Captain Invincible* which is a bit of a cult classic, and remarkable in many ways. I'll spare you the superhero spoof plot, but if I tell you that Christopher Lee performs an amusing cocktail song, you might be interested in checking it out sometime. Except you can't. The video release (PAL VHS version anyway) was deleted some years ago, and numerous rummages in remainder bins and searching on the usual sites will likely prove unproductive. They did for me. The only currently available version is in DVD format. But there's a catch here too. It's region 1 only. There is no region 2 version and there probably never will be.

Now I could just import the region 1 DVD and use some technical trickery to play it on my normally region 2 hardware. But the MPAA doesn't want me to do that, and neither will the EU Copyright Directive (when it comes into force in the UK). The MPAA would like me not to be

able to buy the DVD in the first place probably, but failing that, they'll use copyright laws to make me a criminal for watching it. Does this make sense? Apparently it does to at least half the world's copyright lawyers (guess which side they will be on) and legislators, which is why, delayed or not, directives like the EUCD are changing the relationship between intellectual property and consumers. To find out what's going on and the likely consequences, check out our exhaustive copyright feature starting on page 46.

On a happier note, Mac OSX users can now make use of X-based software. But aside from Apple's efforts to make free software available on the Mac, you should also check out our feature on *Fink* – a porting project that already enables free software greats to run on OS X.

You'll also want to check out two brand-new tutorial series starting this issue. Patrick O'Brien introduces us to the love of his life (Python!) and Michael J Hammel will be showing you how to get some fantastic results out of *The GIMP*.




Nick Veitch EDITOR

The increasingly complicated world of Intellectual Property law and you **p46**

The EPIA M9000 is a pint-sized powerhouse **p23**

Run free software under X on MacOS X with Fink **p54**



LINUX
FORMAT

AIMS OF THE MAGAZINE

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

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

MEET LINUX FORMAT'S TEAM OF WRITERS...



Patrick O'Brien
Published by O'Reilly and IBM and author of *PyCrust* has an enthusiasm for Python that is highly infectious!



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



Richard Drummond
Now firmly in the good old US of A, Rich is on the hunt for a 'decent' beer, with mixed results.



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



Andy Channelle
Busier than a three-legged dog in a forest, Andy compiles the LXF news and writes tutorials and reviews.

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

Bill Von Hagen
Owner of a Lime iBook, our Yank chum has been delighted by *Fink's* ability to run Linux on his one-button wonder.

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

David Harris
Author of our cover feature, David suggests an Open Source Mickey Mouse™ might be the way forward for IP law.

Paul Hudson
The author of our PHP series is currently packing his bags and preparing to move to Bath to join us full-time.

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 p114

LXF38 March 2003

Contents



Welcome to another bursting-at-the-seams issue of *Linux Format*, your guide to all things Linux!



ARTWORK BY DYLAN TEAGUE. JUDGE DREDD IS TM AND © REBELLION, 2003

COVER HIGHLIGHTS

You, Copyright and the law **44**

An Intellectual Property barrister puts UK copyright laws in layman's terms

20 Office without Microsoft

SuSE Office Desktop will lure Windows users

54 X apps on OS X?

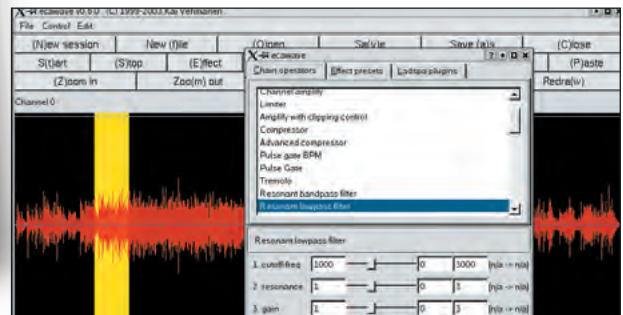
Fink brings Linux to users who prefer one mouse button

72 The GIMP secrets revealed

Professional GIMP user MJ Hammel shares knowledge

80 Learning Python

If you've not scripted before, now's the time to learn!



Which audio tool is best? Judge for yourself in our roundup **p34**

'Spurned and rejected, I thought about giving up programming altogether: it was simply too painful. Then I found Python...'

80 Patrick K O'Brien gets all misty-eyed

'Most of the omissions are sensible, it's easy to bloat a system out including everything bar the GTKitchenSink.'


21 SuSE Office Desktop – Not all there?

LINUXPRO

Linux in your pocket
Montavista CEE 3.0
Grid Computing Does your business need it?
To Protect & Serve A better backup strategy
Tomcat Java servlets & JSP
iptables Set up firewalling from scratch



23 Small = difficult to upgrade, but why would you want to?



6 News

Analysis of the latest Linux stories

12 Mailserver

Got an opinion? Have your say

20 SuSE Office Desktop review

Business apps without the bloat

23 Via Eden review

Good value mini-ITX motherboard

26 Dantz Retrospect 6 review

Back up networked Linux clients

28 Lycoris ProductivityPak review

An alternative to OpenOffice.org?

31 Book reviews

DNS&BIND, Red Hat Bible+more!

34 Audio Editors Roundup

What's best for tweaking your tunes

40 Hot Picks

Pixie Plus gets the **Hottest Pick**

54 Fink

Bringing Linux to Mac OS X

60 Get Help

Where to go other than man pages

62 What on Earth is SVG?

Scalable Vector Graphics explained

68 Mozilla tutorial

Sidebar and throbber manipulation

72 GIMP tutorial **NEW SERIES!**

An introduction to the software

76 PHP tutorial

Creating PDF documents with PHP

80 Python tutorial **NEW SERIES!**

Passionate first steps...

84 Disk partitions tutorial

Every Linux user needs to know!

88 Answers


Your queries answered by experts

95 Back issues

Complete your LXF collection

110 User Groups

Local support, BBLUG's install day



COVERDISCS

DOUBLE-SIDED DVD or 2 CDs
stuffed with Linux apps **98**



CDS A AND B

KDE 3.1 Full release includes KGet, Konqueror, better file management; **Grip** Easy ripping of CD tracks to MP3 or Ogg Vorbis; **phpMyAdmin** powerful yet easy MySQL admin; **Alvaros Messenger** MSN instant messenger network with Linux; **The GIMP** stable and development versions to go with new tutorial; **Qt Parted** Hard disk partition utility; **X Northern Captain** Flexible and powerful X Windows file manager



DOUBLE-SIDED DVD

Cygwin make MS Windows run Unix/Linux tools; **Eclipse** Flexible IDE; **YellowDog** Linux for Apple Power PC users; **Lycoris** Easy to use and maintain distro; **MPlayer** Movie player

Please read the coverdisc instructions on page 107 before installing from coverdiscs!



SAVE MONEY!

SUBSCRIBE TO LINUX FORMAT

See page 96 or phone 0870 4448645



Newsdesk

SCO threatens to upset Linux/Unix/OS X appcart; IBM's Linux growth; New Xandros release; Aqua Data Studio; Bruce Perens books; Lindows Media Centre; Cromwell Linux Xbox mod chip; GnomeMeeting

UPGRADE

KDE 3.1 unleashed!

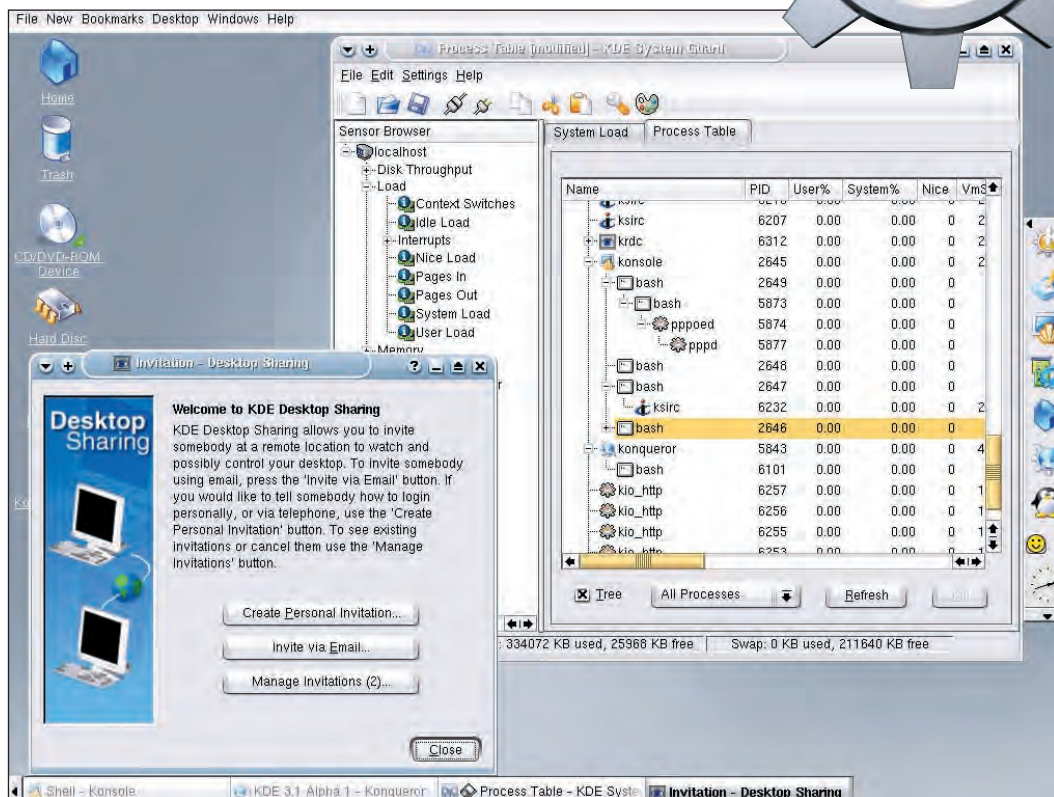


Have this month's lack of major distro releases left you longing for something meaty to install? The final release of KDE3.1 should provide plenty to occupy you. As well as a complete visual overhaul, 3.1 features the highly anticipated upgrade of *Konqueror* (with tabbed browsing), a increased selection of games, *Exchange 2000* compatibility for *KOrganizer's* calendar and *KGet*, a new download manager to simplify file acquisition.

This is nothing less, developers claim, the most advanced and powerful Free desktop for Linux and other Unixes.

KDE Usability Project

UK KDE representative Chris Howells says that the focus on a more sophisticated look and feel is the most obvious change in this version, but there's much more going on besides the adoption of the *Keramik* style and Crystal icons. "Ease of use, and usability in general are hugely important as KDE becomes more popular. A large



The default look is now Keramik and Crystal. But developments in KDE3.1 haven't been restricted to eye-candy.



...and it's on your coverdisc!

number of people work on the KDE Usability Project, and many changes have been made including redesigning the Kicker configuration panel, re-organising the order of modules in *KControl*, and various fixes to other applications such as *Kcalc*."

The most recent LinuxWorld Expo demonstrated the increasingly corporate direction that Linux is headed in, but, Howells says, the support of some of the biggest names in the computer industry has changed little

of the way KDE is developed. "The direction of the project is controlled by the individuals writing the code, not by organisations or committees. Companies largely contribute by donating equipment such as servers, Internet bandwidth, or money to re-imburse developers travelling to conferences."

SuSE's Markus Rex praised the project for its increasing focus on providing a robust and useable environment for Enterprise Desktops. "Features like the Kiosk Mode, the

improved mail encryption and the new Remote Desktop operation are key components of SuSE's upcoming office and enterprise product lines," he said.

The Kiosk Mode has come under the guidance of Waldo Bastian, and provides a simple way for sysadmins to lock down sections of the desktop, allowing more control of what can be changed, installed or removed by end users. The intention was to improve on the standard permissions based systems that have been standard in Unix for years, making

TIMED OUT?

RedHat slashes distro shelf life to 12 months

The biggest name in Linux distros has announced plans to drastically shorten the support life of their basic boxed distributions, a move which,



it simple for Internet Cafes, IT staff and concerned parents, for example, to set levels of access for their customers, staff and children and create "a more controlled environment." The full implementation (with appropriate additions based on user feedback) is expected to be one of the key upgrades for 3.2.

To complement the Kiosk mode, KDE3.1 also has a new VNC-compatible desktop sharing framework intended, the developers say, to improve remote technical support and also to show off your desktop. This is expected to get a major overhaul in 3.2 too.

Linux in government

The enhanced security measures and groupware features have evolved, in part, from the partnership between KDE and the German government which is anxious to implement Open Source software across both local and regional government. The collaboration on the *fgypten* project has resulted in *KMail* acquiring support for important security standards including S/MIME, PGP/MIME and X.509v3 as well as improvements in its OpenPGP implementation. Also within the Personal Information Manager (PIM) application set, the address book library now has support for LDAP and *vCard*, and the calendaring application gets better support for the industry-standard *iCalendar* protocol.

some users claim, makes Microsoft's plans for older versions of Windows look generous. The new 'end of life' policies will effect all 'consumer' OSes, guaranteeing support on the product for 12 months following its release, forcing an annual upgrade on users wishing to stay abreast of the latest security and usability issues. v6.2 and v7.0 will lose errata maintenance support (covering security issues, bug fixes and enhancement alerts) as of March 31, while versions between 7.1 and 8.0 will be dropped at the end of 2003. Conversely, the Red Hat Advanced Server is given a shelf life of three years, prompting speculation that the company is hoping to bounce more users onto the higher cost solution.

Forward Planning

Mark Cox said that in the past errata support would usually be discontinued when a major upgrade was released, and the new schedule would allow Red Hat to concentrate on ensuring timely support for current releases. "We believe Red Hat Linux users should be able to plan migrations and upgrades in advance and therefore need an errata support policy that gives clear guidance on product end-of-life dates, without any immediate discontinuation as we release new major revisions," he explained, somewhat tortuously.

Posts on RedHat's announcement mailing list were split on the issue, but a slim majority thought it was a bad decision: "From a business point of view, one reason I do buy Red Hat CDs is to make sure Red Hat is viable in order that it can continue to provide support. It's not the only reason, but it's an important one, and if RH lessens its support, it means that reason is less of a reason now," wrote one disappointed user. Another suggested reducing support was probably in Red Hat's best interests: "Support takes resources, but it doesn't generate any revenue other than good will. It's likely that the drop in good will be less than the savings in support costs." There's been no comment so far from IDG Books Worldwide, publishers of the *Linux For Dummies* series, who ship their popular range of books with Red Hat CDs.

NEWSBYTES



■ Proving that time waits for no-one, supercomputer prime-mover **Cray** is attempting to develop a successor to the X1, even though it has only been available for a short time. Codenamed *Black Widow*, the new machine should eclipse the 52.4 Teraflops of the X1 by a large margin. A spokesman mentioned power in the range of 'hundreds of teraflops'. The problem Cray have is that due to potential licensing issues with SGI and Irix, the company may have to fish around for a new OS for the X1 offshoot. Have they considered Linux?

■ **Sun** have become the first member to release a Liberty Alliance compliant 'identity server'. Version 6.0 of SunONE Identity Server uses version 1.0 of the Liberty Alliance specification which was finalised in July 2001, but a patch will be issued for the 1.1 spec due in the very near future. The Liberty Alliance is a system designed to provide single, federated sign-on facilities for compatible accounts and is intended to compete directly with the ubiquitous *Microsoft Passport*.

■ **MontaVista's** embedded Linux distribution has gained instant messaging facilities thanks to a deal with Movial. See this month's Linux Pro supplement for more information on this development.

■ Is this another signalling of the warming relationship between **MS** and Linux, a more pragmatic approach to a growing market, or a shot across the bows of Real's *Helix* project? MS has added a little flexibility to its licensing terms for the DRM technology in *Windows Media Player* (9 series) that, potentially, allows the development of a Linux version of the application. It was only six months ago that various MS executives were claiming that such a move would mean losing control of their intellectual property...

■ **Goldman Sachs** say 'Fear the Penguin'. In a recent report, the investment bank claimed that Linux would make sweeping gains across various sectors at the expense of both Windows and proprietary Unix. The task for Linux-advocates, the report says, is to refute the impression that the operating system isn't 'Enterprise Ready'. 'Linux has evolved into a enterprise-class operating system', the report stated. 'But most corporations still appear to view it as nascent technology.'

David Cartwright

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



COMMENT

Real-time Java?

“ Everywhere I go, people extol the virtues of Java – how it lets you do stuff quickly that wouldn't be possible in languages like C, and so on – for free. Though I've never escaped the clutches of C, ASP and VB, I can see it has a lot of benefits. I'm wondering: is Java any good as a server-side Web language?

Having worked in several largish companies as both employee and consultant, observation tells me that 100% of those with Java-based websites have very slow sites. Am I missing something? Is everybody I know with a Java website simply incapable of configuring a server correctly? Or is it just that Java's so versatile that it compromises on speed?

I know two companies with different Web approaches. One has a simple Perl-based system reading its data from text files when a search takes place. The other has a big, professionally built Java-based system back-ended with *Oracle*. The former is significantly faster than the latter, and has 80% less hardware too. In another client, we chucked away Java, which had been used for the first chunk of an internal project, and used ASP instead and everything went an order of magnitude faster with the same *Oracle* back end database.

Might this be something wrong with the database connectivity aspect of life? I can't help thinking it could actually be that Java is too general for it to be expected to stand up to a pounding on a site that's trying to do quick searches for its customers? Please email linuxformat@futurenet.co.uk – I'm interested to hear readers' experiences of Java in sites that try to be real-time. ”

BOISE WILL BE BOISE!

SCO to seek UNIX patents from Linux users?

Reports indicate that SCO (formerly Caldera) will attempt to enforce patents it owns on the UNIX operating system, a move that could directly impact on Linux use and development. The company has hired the services of attorney David Boise (who prosecuted Microsoft and defended Napster) to investigate whether Linux, MS Windows BSD or Apple's OS-X infringe on patents granted in relation to SCO UNIX. Earlier SCO had announced a plan to probe the use of two patents it owns



on libraries which allow SCO software to run unaltered on Linux, but this 'widening of the net' could, according to some analysts, have disastrous consequences for SCO's Linux strategy. Stacy Quandt of the Giga Information Group said in an interview that SCO

risked putting themselves at odds with large numbers of potential customers.

Darl McBride, SCO Chief Executive, said the investigation was not a recent strategy. "We've been looking at this for months," he said. "Every time we turn over a stone, there's something there. If

you pull down (Mac) OS X you'll see a lot of copyright postings that point back to Unix Systems Laboratories, which is what we hold." All this legal wrangling could have serious knock-on effects for the whole Linux world, we'll endeavour to keep you updated.


Linux Web Watch/



Cutting edge development.



Dispatches from the legal frontline.



jwz's uncompromising front page.



Domestic strife = book + movie!

Blogging a dead horse...

Is Weblogging merely 'vanity publishing for the digitally aware', or are there advantages to be had from maintaining such a site? LXF searches for content of consequence amid the typo-strewn and badly punctuated ramblings of the majority.

For all the talk of a democratising force, the Internet – in the main – has simply provided another mode of advertising and reaching a serious number of hearts and minds all too often depends on blowing some serious cash. Web Logging (blogging), seems like a return to the text-based roots of the Internet where content is king. Linux users, tech-savvy to a man, are quite common in Blogland.

Mitchell Kapor (<http://blogs.osafoundation.org/mitch/>) was one of the team responsible for Lotus 1-2-3, the original Killer App that demonstrated a simple way that

computers could have an enormous impact on life and commerce. Recently he's channelling his energy (and some serious financial muscle) into Chandler, an Open Source all-encompassing communication tool. Most of those involved in the project are committed bloggers, but the most interesting stuff eventually finds itself here.

In the battle for rational copyright law in the USA, Lawrence Lessig (<http://cyberlaw.stanford.edu/lessig/blog/>) is a tireless campaigner. His well-designed and active weblog details the many efforts he's made, including the recent (failed) challenge to the *Sony*

Bono Copyright Term Extension Act, to curb powers of content owners. Lessig is also chair of the Creative Commons (www.creativecommons.org) project.

After bailing on the *Mozilla* project Jamie Zawinsky (www.jwz.org) set up a weblog-like site from which to hurl random tracts of venom at various topics. Despite occasionally coming across as Victor Meldrew with too much black clothing, Zawinsky is a witty writer and gets to the nub of his argument with little in the way of bush-beating. Lots of swearing too.

Not strictly related to computers, Linux or open source, but possibly one

of the funniest blog sites you'll find – www.thingsmygirlfriendandihavearguedabout.com is essentially an account of the various arguments Mil Millington (who started his writing career on *Amiga Power*) has had with Margret, his "dangerously insane" German girlfriend of twelve years. Proof that blogging can lead to greater things than a diminished RL social circle, Mil has just finished a serialisation of the site in *The Guardian*, and also published a book and acquired a movie deal, both based on a fictionalised retelling of his blog events. See also the quaint and amusing www.theweekly.co.uk

NEWSBYTES




■ **Xandros** are releasing a 'Standard Edition' of their Debian-based operating system which removes *CrossOver Office* and knocks \$60 off the price. The new version also loses the ability to reassign space from NTFS partitions and drops the 30 day installation support in favour of a single incident email option.
www.xandros.com

■ Hoping to help speed up development of the next generation of desktop environment, some bright spark at **KDE** has come up with an 'adopt-a-geek' scheme. "Many of these folks (KDE hackers) overwork their underpowered computing setups to help bring you an amazing Open Source desktop". Citing a compile time of six to eight hours for the latest version of the desktop on even a decent PC, the project aims to unite needy developers with generous end users who may have too much hardware on their hands.
<http://devel-home.kde.org/~wheeler/adopt-a-geek/>

■ Unilever, one of the world largest commercial organisations and owner of many popular High Street brands, has shifted much of its computer needs to Linux on the basis of cost, power and stability.

■ Investment Bank **Morgan Stanley** has switch one third of its server hardware to Linux in a bid to reduce IT budgets. A spokesman said this was the first step in moving the whole of their IT infrastructure to commodity hardware.



■ First-Person Shooter and recruiting method for the US Army, **America's Army** should soon have a stable client for Linux. Despite its flag-waving premise, AA, based on the *Unreal Tournament* engine, is notable for allowing both sides (you play against live players on the Net) to be the good guys.

BILLION-DOLLAR REVENUE

IBM reveals strong Linux growth

IBM's Head of Software, Steve

Mills, used his keynote address at New York's LinuxWorld conference to reveal how Big Blue has managed to make money from Free Software, and to predict which way the market will turn after 'the hardware rush' is over.

Predicting annual growth rates of up to 35 per cent across the industry, Mills said that Middleware and services were taking over from hardware as core revenue earners, especially in IBM's target markets – data centres, distributed computing environments and retail. Mills also confirmed that Linux was the 'logical successor' to AIX, but that the transition would be extremely long-term, i.e. decades rather than years.

This year, IBM also shifted their emphasis. Previous keynotes had concentrated on the company's intention to make Linux ready for the big time, most notably the '\$1 billion



Steve Mills used his LinuxWorld address to highlight a number of major organisations buying into Linux and IBM.

development program' of two years ago, but this time new General Manager of Linux at IBM Jim Stallings announced that Linux software, hardware and services sales topped US\$1 billion in 2002.

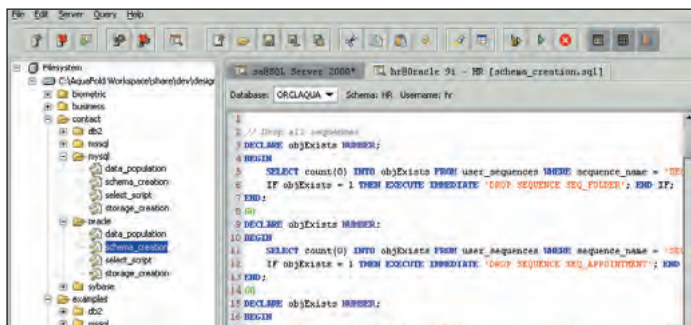
DATABASES

Free-to-use database development environment

Aqua Data Studio (ADS) is a free (though not free in the FSF sense) development tool which provides facilities to browse database structures and create, edit and execute database scripts. With support for any Java enabled operating system, ADS provides a single consistent user interface to many relational databases including Oracle 8i/9i, DB2 7.2/8.1, Microsoft SQL Server 7.0/2000,

Sybase ASE 12.5, PostgreSQL and MySQL. The product's developer Aquafold claim ADS's sophisticated graphical browser gives users a convenient overview of the structure and dependencies of a database with the click of a mouse. It also feature auto-completion of SQL statements and the ability to output query results as HTML, XML or CSV.

www.aquafold.com



ADS can browse relational databases from a diverse range of vendors.

Jono Bacon

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



COMMENT

Limiting the innovation?

“ One of the overwhelming benefits of the Linux and Open Source world is that anyone can contribute. Given this freedom to innovate, many do get involved. Although this innovation is greatly beneficial to get an application off the ground, is there a point where the innovation should cease due to usability coherence?

'Bloat' is common in competing products, and one that stuffs many proprietary applications with pointless features that are rarely of use. With Linux and Open Source software gaining a degree of feature maturity, there is risk that this same issue could be a problem to the popular Open Source apps we know and love. How do we place a limit on the feature innovation while not putting off the developers who eagerly code into the night for these projects?

It is a fine line, one that must be tread with a degree of modularity, so unwanted features can be easily stripped out without misshaping the general usage of a product. It may be worth considering for the stability, efficiency and size of the apps whether that new icon set is worth the extra bloat. Features are without doubt an important factor in a product, but coherence of the user experience is more important, the 'feel' of an application greatly impacts its success out in the field.

I have no doubt that Linux and Open Source hackers will find suitable resolution to the problem, but bear these factors in mind next time you email a developer asking for that killer feature. It is the responsibility of us all to help keep a balance of features and efficiency level, even if this warrants holding back on some of the innovation. ”

FAMILY PC

Lindows Media Centre

Lindows.com are expanding their operations into the hardware arena with the Lindows Media Centre, a small-form PC designed to serve up DVD video and music (CD and MP3), while still allowing users the full advantage of owning an Internet connected (LindowsOS) PC.

There are a number of features, not least the price, which elevates the Lindows Media Centre above the norm though. Firstly the size and style suggest a product designed for the living room while the price, less than \$350, is designed to encourage us to buy a couple (Michael Roberts, Lindows Founder and CEO says he envisages a future where each house and office has a number of these things – he's not short on ambition) of the things for the home network.

Beyond these two factors, the most interesting thing about the LMC is its

hardware. While the processor (VIA e-series 933MHz), memory (128MB), harddisk (20GB) and everything else is pretty much at the lower end of bog standard, the device features Elegant's www.elegant.com etDVD BIOS which allows for the playback of DVD movies without bothering the actual operating system. As well as reducing system overheads, etDVD also gets around problems with decoders, CSS and the non-intuitive nature of the current crop of Linux DVD players.

So, for the cost of a decent DVD player, users get a pretty good



PC/media server thrown in too. This could be the niche that Lindows makes its own.

Small, stylish and capable: Lindows Media Centre should turn out to be a market-leading product.

BRUCE PERENS BOOKS


Free Treeware

Free software pioneer Bruce Perens has given his name to a series of books to be published under Open Source licenses, a move which Perens calls a "tremendous departure

for a mainstream publisher". Prentice Hall have already published two books, and it is expected that the series will feature between seven and nine books by the end of 2003.


Perens said the texts were released under the Open Publications License 1.0 with no options taken, meaning they can be copied, modified or redistributed freely without additional permission. A few months after their traditional release, the books – which will be primarily concerned with Linux and Open Source software – will make their way to the Internet.

"I swear I didn't ask them to name [the series] after me, it was their idea," Perens says. "They need my name there to differentiate their line from any other Open Source book series that another publisher creates. I will go along with



Embedded Linux News

● **e-LAP** is a new PDA reference design created by IBM to jumpstart the Linux PDA business. The design can be taken and adapted to form a range of devices from standard PDAs to fully-featured persona media players. The platform specifies a PowerPC 405LP Processor, 32MB SDRAM, and the controversial TCPA security chip. Software includes MontaVista Linux Consumer Electronics Edition (Trolltech Qt/Embedded, Qtopia PDA application suite, and Opera browser.



● **American Arium** have developed a new kernel level debugging toolkit aimed at coders working on ARM and XScale processors. Todd Selbo, product marketing manager said the toolkit would allow developers to see their sourcecode executed on a line-by-line basis.

● A porting project to bring Real Networks' Open Source Helix DNA client to the Zaurus has been initiated. Once the Zaurus has been conquered, the project

team aim to widen their ambitions to include all ARM-based devices including set-top boxes, mobile phones and PDAs. <https://linuxarm.helixcommunity.org/>

● A new **Xbox mod-chip** from OzXchip has begun shipping complete with the Open Source Cromwell Linux BIOS and a boot CD called *OzXFlash* which enables users to flash a new BIOS onto the mod chip simply by boot up with the CD in the drive.

that, as long as they are willing to help me get good documentation into Open Source."

The first three books to hit the shelves are *Embedded Software Development with eCos*, by Anthony

Massa; *The Linux Development Platform*, by Christopher Paul and Rafeeq Rehman; and *Intrusion Detection with SNORT*, by Rafeeq Rehman. Visit <http://perens.com/Books> and <http://vig.prenhall.com> for more.

ALTIX SGI's very big iron

Silicon Graphics Inc (SGI) started its new year with a bang by realising the ambition of bringing a 64-way Linux/Itanium Supercomputer to market. With NUMAFlex clustering technology adapted from the Origin 3000 range, SGI hopes the Altix line will finally allow them to compete on the same territory as Dell, Hewlett Packard and IBM, while bringing speed and stability increases to the sector in general. NUMAFlex offers substantial improvements in memory management, which SGI claims is up to 200 times faster than the ubiquitous Myrinet switches used in the average Linux cluster. An SGI spokesperson said the company would be targeting existing users of Intel-based Linux clusters who could 'drop their applications right into an Altix3x system unchanged and see a dramatic improvement in performance.'

There are currently two versions of the Altix available: the 3300 is built

around four to 12 Itanium2 processors, while the 3700 scales from four to 64 processors with upto 512MB RAM per node.

And if that's not enough, SGI are promising a super Supercluster which bands together eight 3700s and provide one terabyte of shared memory. The system, SGI claim, can scale up to 2,048 processors without changes to the Linux kernel.

The Altix won *Best of Show* at the recent LinuxWorld Expo.

The new Altix line boasts insanely powerful hardware and an eye-watering cost.



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

It's time to stop

“Buggy software is so common that it is assumed that software can't exist without bugs. Just look at the Bugzilla lists of the big projects and examine the number and age of outstanding bugs. Some niggling bugs have survived several major version releases. Why?”

It seems as if we have to choose between adding a new function or fixing a small bug, the bug-fix loses most every time; only the major show-stopper bugs get fixed. The common rebuttal is to complain that bug-free software can't ever be shipped since features drive development and features are what the buying public demands. So how does this thinking fit into Free Software?

Free Software *should* be different. While distro makers seem driven to release by the calendar, the software developers should not be. What I propose is a moratorium, a pause, a recurring moment of sanity. We need a Linux-wide feature-freeze for one month a year that is devoted exclusively to bug-fixing. The benefits are many. Fewer bugs means fewer problems, better reliability, and easier acceptance from new users. In general, kernel development appears to be maturing and slowing its pace. Let's use that to improve software quality. I propose July because it comes before the August holidays and it prepares everyone for the calendar-driven distro releases of software for Santa-day.

KDE recently delayed the release of version 3.1 for security fixes. I hope they got a lot of bug-fixes in as well. Let's make regular bug-fixing a part of the Linux culture.

CHEAP INTERNATIONAL CALLS GnomeMeeting PC-to-PC phone hardware solution



GnomeMeeting's new dialler both looks elegant and does the job.

GnomeMeeting has taken another step towards becoming the ultimate contact utility, gaining VoIP facilities and dialling of standard phones from the desktop. Users must first purchase Quicknet www.linuxjack.com hardware – a Linux-specific soundcard with telephony facilities with prices starting at \$109 – and an account with MicroTelco www.microtelco.com.



Linux can now make calls to phones, after you've got relevant hardware.

Once you're set up, you can make free PC-to-PC calls and fairly cheap overseas calls. The latest release (0.96) also features a new configuration / testing druid for audio and video, drag-and-drop support on both the address book and ILS browser and full internationalisation support in H.323 and ILS protocols. www.gnomemeeting.org

CROSS-PLATFORM MAIL Mulberry Mail Client

Cyrusoft have announced version 3 of their *Mulberry 3* cross-platform mail client for both Linux and Solaris. The new version is a little less eccentric than previous iterations and now follows the standard three-pane look favoured by *Evolution*, *Mozilla* and everyone else. Separate mailbox and address windows have a preview pane, and click-on-column sorting is also supported. Configurable actions make personalising your options easier.

CEO John Degory announced the release at New York's LinuxWorld Expo, saying sales of the mail client had been growing steadily within the Linux community over the past year. "this new release will allow this growth to accelerate," he said.

A new tabbed interface provides easy access to multiple mail accounts, and a macro feature makes inputting commonly used phrases and text sections simple and quick.

Mulberry 3.0.1 is available now, visit www.cyrusoft.com/mulberry

Mailserver

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

All the Same?

I just want to thank you for a very eye-opening view of Linux with your *DemoLinux CD*. I must admit that I am a Windows user, and know extremely little about Linux except that it is said to be much faster and more stable than Bill's Monster, but very user unfriendly, a b**** to set up and not easy to find software for. I have noticed on the web that the majority of programs/websites appear to be done with Linux and it seems to be a very large, though confusing, community out there for support. After suffering from Bill's BSOD one time too many, your mag caught my eye with the "run off CD" and I am very impressed with my initial experience though I am really lost in this different OS with so many strange-sounding words. I want to learn more about it, though I'm not sure if I'll be able to grasp the whole picture. So it seems that Unix, Red Hat, BeOs, Os2 are all open source Linux in different forms?

Peter J.A.Hall, South Africa

Glad you have had the courage to try out Linux given your preconceptions. The short answer to your question is that only Linux is Linux. The other operating systems you mention only have the fact that they are not Microsoft in common. There are various flavours of Linux, names like SuSe, Red Hat, Mandrake, Lycoris, Xandros and others. These are all variants on a common theme but share the fact that they are all based on Open Source software and use the Linux Kernel as the core.

There are other Open Source and free operating systems which are similar to Linux, but which do not use the Linux Kernel. An example here would be FreeBSD. It is similar to Linux because it shares many of the same applications, and has a UNIX-style environment, but as it doesn't use the

Linux Kernel, it isn't Linux. There are also many shades of Unix, most of which are proprietary. Although they generally support the same software, these often use closed code too. Examples would be Sun's Solaris and HP's HP-UX versions of Unix.

BeOS was originally conceived as a commercial alternative to Windows, but never really took off. The original company hit financial trouble, but the project has been resurrected as OpenBeOS, an Open Source, free, rewrite of the system.

I hope that has given you something to think about - not everything that isn't Windows is Linux!

Chain letter

I recently decided to upgrade my computer and went to some of the big (New Zealand) Chain-stores in search of a new system.

As a Linux user of about a year's duration, and I consider myself as a bit more than a newbie. I made the conscious decision that I wanted to have my new machine as a Linux only box and get away from Microsoft totally. In each store I went to, the sales person involved was quite definite in their response "You don't want Linux, Windows is the only way to go!"

In fact, in one store, to get a computer to the spec I wanted, without an operating system, was much more expensive than a higher spec computer with Microsoft XP and their large software bundle; and of course with XP pre-installed, one can't un-install it and sell the disk.

Another store I went to, the first salesperson tried very hard to talk me out of setting up a Linux box, and went on to show me a vast

array of computers already installed with XP. After the sales person saw I was really not interested in what he was trying to sell me, he walked off. I was having a final look around the store when the manager of the computer department came over to help me. I stated that they obviously couldn't help me with what I wanted, and was just about to leave.

Well, to cut a long story short, they eventually did sort me out with the system that I wanted, without XP, and threw in a DVD/CD-RW combo drive instead. So, I now have a system running Mandrake 9.0. All I need now are stickers on my box pledging my allegiance to Linux!

The one thing that annoyed me in every store I went to, was the fact that most salespeople just assumed that I didn't know what I was doing. The whole experience would have been more bearable if the salespeople actually *asked* if I knew what I was doing and not assume me to be some ill-educated idiot (which was the way I felt I was treated in most of these stores).

I know I could have avoided the episode by going directly to a specialist computer shop, but they weren't doing the '24 months interest-free, no down payment' deals, which, I am sad to say, was the deciding factor for the whole humiliating experience of going to these chain-stores.

On a lighter note, Thanks for the great mag, I enjoy getting it every month, even if it is 3 months late over here in NZ. The CDs are very comprehensive and I love trying out the different distros.

Rob Martin, Christchurch, New Zealand
Your experience is repeated the world over. There is no commercial incentive for stores to sell Linux, and therefore, they rarely learn anything about it. In fact, we're very surprised the





salesperson had even heard of it! The general attitude is the same most places too. I guess we are all susceptible to stereotyping though – I know I always automatically assume I know more than the salesperson, (er, though most of the time this appears to be true). Well done on a positive result though!

Working Weekend

Thanks, LXF! You just saved me from a fate worse than ... well, perhaps I exaggerate!

At the weekend of 01/12/02, I thought I would try Linux. I had enough discs from various mags, and I chose a DVD with three images to burn Mandrake 8.2 disks. Along with another disc with a copy of Partition Magic 5, I created a new partition, just 6GB, for Linux. PM5 also “advised” me to create a swap partition. I didn’t know what it was, but I went along.

Then I loaded Mandrake. Easy-peasy... or so I thought! It all seemed to have gone well, but at the end of it I was presented with a 800x600 screen (not the best when you have a 1280x1024 LCD monitor), no sound, no printer, and no Internet.

Running Windows 98SE, I thought I would search the Internet for advice. It turns out that my Motherboard, a PC-Chips 810LMR, has SiS chipsets for just about



Mandrake 9.0 – straightforward even in ‘expert’ mode.

everything (SiS630 for video), and that they will not give out the info which Linux coders use to write the necessary drivers. I tried posting to www.linuxnewbie.org and found a particularly depressing site which virtually wrote off SiS hardware as far as Linux was concerned.

The only positive was that Mandrake 9.0 was out, and there may be improved support for difficult hardware. I tried to download the disc images from the Net (ntl 128k broadband) but found in the morning that the download had been broken. I was fed up and on the verge of giving up... Then, on the way home from work, I spotted LXF on the shelf in WHS, with Mandrake 9.0 discs attached to the front cover. Looking at your installation instructions, I even chose “expert” mode and found it quite straightforward.

I now have sound, my ntl Broadband works after a fashion, I can get my monitor in max res at 16-bit colours (and I cannot tell the difference from 32-bit). I still have a few problems. When I am offered 1280x1024 24-bit graphics by the

(Connected by a USB Cable Modem 351000 on eth1 using dchp). The most frustrating thing is that nothing I have tried (and believe me I’ve been to *all* the sites and got *all* the software) gets my printer to work. (A Lexmark Z54. The Z52, Z53 and Z55 all work, but not mine, Argh!)

Still, I’ve made enough progress to want to persist. Had your cover disk not made any difference, that would have been me and Linux finished. Now I will be running a dual-boot system with Win98SE, and hope that, by the time pre-XP Windows becomes obsolete (as I understand that new programs may not support it), that I’ve learnt enough about Linux, and there is enough software out there for me to

“There is no commercial incentive for stores to sell Linux, and consequently they are unable to offer prospective users any advice about using it.”

HardDrake, and select it, I get 1024x768! When ntl dynamically update my IP address, I lose my Internet connection and HardDrake cannot reconnect it, only restarting the computer seems to work.

face the future without Windows.

If anyone knows the hardware items I mentioned, maybe they can advise me?

Dave Spagnol, *via email*

Glad we could be of service. You are



★ Letter of the month

This month's winner receives a copy of **RedHat Linux Internet Server**

Geektivism

There used to be a saying that ‘nobody ever got fired for buying IBM’. Today’s equivalent appears to apply to Microsoft. Despite ever-deteriorating products, where the only new ‘innovations’ are geared towards trapping customers into an ever-more expensive upgrade trap, most managers appear to have no more sense than to take the immediate path of least resistance.

So many times, in workplace situations, I have encountered ‘problems’ to which Linux solutions are readily available, but man-hours are squandered cobbling together fragile and expensive Windows-based solutions, and where this is not possible (*ie* very often), the problem

is deemed to have ‘no solution’. For example, I’ve seen an instance where standardisation of applications and desktops across an organisation was required, could easily be implemented using ‘apt’, but was deemed ‘impossible’. I’ve also seen a roll-out of NT with reflections for a large number of PCs that simply needed a ‘telnet’ session. How painful does it get?

However foolishly organisations squander their own resources, it’s still ultimately their own business. It becomes everybody’s business when these stupid decisions cause low-quality servers to pollute everybody’s internet, and cause the chaos that has recently been seen by Code-red, Nimbda and most recently Slammer/Sapphire. It should now be

obvious, that hooking such servers to the Net, however patched, lies somewhere within a triangle between stupid, irresponsible and anti-social.

I’m open to suggestions as to how to contribute best to stamp out bad net practices. Maybe it would be practical to do some sort of *Roundup* section on organisations such as Anybrowser.org, Campaign for Digital Rights, Insecure.org and other such sites, to quickly outline their motivation and objectives. Then, those of us who care about our Internet can constructively act to protect it.

Mel McWeeney, *Ireland*

Poorly maintained servers are a painful burden for everyone else on the Internet, there is no doubt of that. You’d be surprised how many ‘Code



Red’ attacks we still receive on the *Linux Format* webserver for example. Although harmless, they still clog up the logs and waste bandwidth. I would hope that anyone encountering a badly behaving server would contact the organisation concerned (and we’re sure there are poorly behaved Linux servers as well). Thanks for the suggestion, I’m sure we can cover this in a future issue.

And to make sure you are well versed in keeping your server up to date, we’re going to send you this month’s star prize, the *Red Hat Linux Internet Server* book, the official Red Hat guide to running Linux on servers (ISBN:0-7645-4788-7) written by Paul Sery and Jay Beale and published by Red Hat Press/Wiley.

Mailserver

« not alone with your ntl problems. Many Linux users seem to have had the same problem. I thought the issue was related to the 'pump' DHCP client used by some distros, but as Mandrake 9 doesn't use this, there may be some other problem. There is an ntl newsgroup for Linux though, so when you do get online you may find some useful advice there.

Where's mICQ?

I work on the Japanese translation for *mICQ* <http://micq.org/>. *mICQ* is a GPL licenced terminal client. Please would your team consider reviewing this ICQ client as it was not included in your original Instant Messenger review? *mICQ* supports the new v8 (AOL OSCAR) protocol more

"An enormous problem for newbies switching over from Windows is making an Internet connection – and Linux isn't desktop-ready until this is fixed."

completely than *any* other Free Software ICQ client. It is convenient to use because it is terminal based; I can be on ICQ from anywhere, I just SSH into my machine. The key feature which I selected it for is its use of Unicode UTF-8, this means I can communicate in any language. It has other advanced features like identifying other clients.

Don't be put off by the low version number, *mICQ* does not conform to the standard 1.0 is first major release ideals. *mICQ* is very feature complete, and beats other clients on virtually all issues.

Regarding the legal battles you mentioned with the free IM clients that are compatible with AOL OSCAR protocol servers, I have not heard about these. Officially AOL has to open up their protocol when they integrate 'next gen' services; this was part of the agreement that let them merge with TimeWarner. Legal battles relating to their use of trademarks are expected, but being compatible with their protocol does not seem to be an issue for them. AOL are much better than many, AOL funds *Mozilla* which is a project covered by GPL (and other licences). JGrant, via email

Thanks for getting in touch. As I'm sure you can appreciate, very often in the roundups we can't include absolutely everything, though we do try to include the most active or widely used pieces of software in a category. Sorry we missed out *mICQ*, but we'll take a look and maybe feature it shortly on its own.

Not Ready!

I decided to see if Linux was something for my home desktop when I got disillusioned with MS Windows Me just over a year ago, so your article on this subject was my number one read in the February issue. Congratulations to Andy Channelle for the in-depth look at the factors involved. The Web comes



Maybe our question should've been "Is the desktop ready for Linux?"

with the Mandrake distribution on your cover discs. This problem is a serious stumbling block and until the Linux community really addresses it, Linux is not ready for the desktop as far as I am concerned. The only other option is to buy a Linux-compatible modem which I refuse to do; it would be interesting to know how Evesham have solved this problem in their recently launched Linux-ready low-cost machine.

This is not the only shortcoming; the ones that trouble me are firstly, although there is 'help' on most applications it is frequently inadequate and difficult to understand, also it varies too much in quality between applications. Secondly the fonts displayed on KDE seem to be fuzzy and consequently off-putting.

I would add *Media Player* and TV card to your list of the six essentials for the home user. Again there is a problem with my *Hauppauge WinTV*. Although Mandrake include the

TVcard in their installation, it won't work and *Hauppauge* will only support Red Hat and SUSE. So I'm stuck again.

Glad to say it's not all hassle though, as I find the *Gnumeric* spreadsheet superior to the one that came with my PC (*Works Suite 2001*).

I believe your article will have done a lot of good towards achieving the object of getting Linux firmly established on the desktop if it receives plenty of letters, such as this one, describing users' experiences so leading the Linux community to drawing up a list of its present weaknesses and having a blitz at fixing them all.

Thanks for an interesting and informative magazine.

S. Ionii, via email

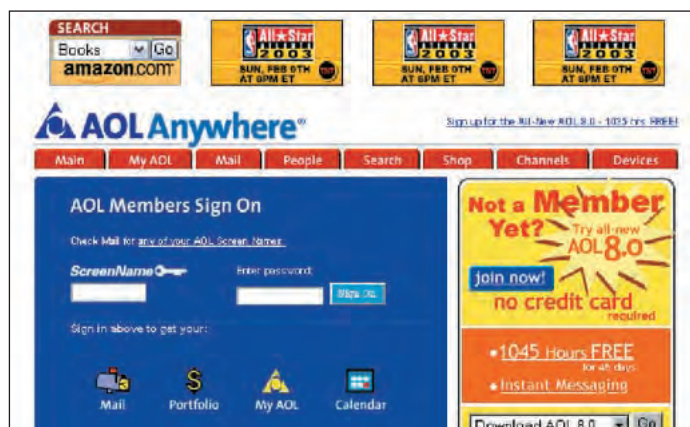
The problems with Winmodems have been around for a while now, but it isn't one that the Linux developers can easily solve. The problem is that these devices talk directly to the CPU, using its processing power to complete some tasks which would otherwise require additional (equals expensive) hardware. Unfortunately, there is no real standard for these devices, so they tend to behave in different ways. If open source developers knew exactly what the Windows drivers were doing, it would be possible to create Linux drivers for these devices. The majority of the manufacturers do not wish to share their code, or even describe how the drivers are working. Sadly, this is not a problem that can be solved just by getting more developers working on it.

A few companies, for example Lucent, have helped create Linux drivers for a variety of devices using their chipsets, but it's a small drop in a big ocean.

Documentation is something that can be solved though, and quite easily. There are plenty of people using Linux who think they can't contribute to free software because they have no programming skills. Well, most people can write! If you find something with poor documentation, contact the developers and offer to help them with it!

Omission

Page 71 and 73 of *LXF* 34 contain potential errors (that can be serious in a 'real' system) - an error that had me going round in circles many



mICQ – would TimeWarnerAOL's demerger affect source code availability?



years ago, wasting a weekend. I remember that I used:

```
truss -p <pid>
```

(on Solaris) to discover the problem. Unfortunately though, some Linux variations don't have this command.

Anyway, here's the problem: As stated in your tutorial, the **accept(...)** command blocks and waits for a client connection. But, the code doesn't consider an interrupt occurring causing the accept to return prematurely (On Solaris this used to happen a lot and Linux systems need the same check).

If **accept()** fails, returning -1, your example programs still go on to attempt to read via the returned socket descriptor (e.g. -1) and then continues processing the data expected to have been received, and so on. The return value has not been checked to ensure that an actual socket descriptor was returned.

In summary, if an error occurs, it is worth checking the global **errno** variable (may need to include file **error.h**) in the case of **accept()**. Here is an example:

```
if ( (fdnet = accept(sock,NULL,
NULL)) < 0 )
{
    if (errno != EINTR)
```



If you had problems with LXF 34 *What On Earth*, Elaine's got the solution.

```
{
    perror("cannot accept
requests");
    exit(1);
}
....
```

With systems programming, it's important to emphasise error checking like this. I'm so glad my students didn't see this tutorial. Their recent assignment was to write a webserver. I know your example was probably kept simple to make learning the topic easier. Elaine Wright, College Tutor
elaine.wright@ntlworld.com

Thanks for the correction. As the example was showing the use of **accept()**, and space was tight, some of the usual error checking you might perform in a real environment was removed. But thanks for highlighting this issue! I hope your students wouldn't be copying their homework from LXF too...

JigDo alternatives

Since I was the one complaining about JigDo, let me suggest a couple of alternatives. Whether they're practical or not is another matter, but never let it be said that I haven't tried.

1 Just use a Perl program to re-create the images. For example, there are lists of which files go where for the Debian install (I found them on the Net). A simple perl prog that takes a dir as an argument will allow copying things to that, sub-directoryed into CD1, CD2 etc., and will let me create the iso images from those dirs. Most UNIX-based boxes will have Perl installed and it runs on more things than JigDo does.


2 Take a leaf from the floppy boot concept. My Internet gateway can't boot from CD. I have to boot from floppy and then tell the installation procedure to read the CD.

If you can get the floppy boot procedure onto the DVD then it should be possible (stress: *possible*) to tell that procedure to look at a bunch of CD images (or mounted CD images) to continue the process. Whether this is doable for every (or indeed any) distros is another matter.

Yes, I know what the purists will say but then if you ask 15 experts how to do something, you'll get 16 answers. In the meantime, I've run through the Linux From Scratch install on the latest CD. This is neat! I greatly appreciate you putting it on the DVD. I've learnt a number of things >>

Helpdex

shane_collinge@yahoo.com



Mailserver

« already and I have a greater respect for the distro makers now. Actually, it almost makes me feel like a real sysprog again!

Rod Furey, *The Netherlands*

You can already use a floppy image to install from, although generally to work this will require two machines networked together.

Both of these are probably workable solutions, but would require so much effort to set up for each distribution, that they simply wouldn't be possible. Both are going to be specific to the distribution included, and we don't have the time, never mind the manpower, to generate and test scripts and options like this on an issue by issue basis.

I agree that JigDo isn't perfect, but it's the best we have at the moment. If you can think of a way to build a generic script for this that works, or if anyone has further ideas, please let us know at the usual address!

Simple mistake

With reference to the letter *Simply Linux (LXF36)* concerning a salesman who thought it was illegal to sell a PC without MS Windows, I will explain the typical scenario for

“We wouldn't buy new copies of any other software we already owned, so why should we have to buy an OS with each new unit of hardware?”

a computer sales company, let's call them *Cheap PC Bloke*.

Cheap PC Bloke sells a large volume of PCs and they require an Operating System for their products. *Cheap PC Bloke* looks at the market and chooses an OS based on it's popularity; let's call the OS manufacturer *Crash And Burn (CAB)*.

Cheap PC Bloke is looking for an edge in the market and explains to *CAB* that they wish to purchase a large volume of *CAB*'s product and are looking for a volume purchase discount. This could considerably reduce their production costs and allow them to undercut their competitors, thereby increasing their market share.



Could porting be the answer to a harmonious closed/open relationship?

Now *CAB*, being a bit of a sly fox, says that they will supply their products at a considerable discount as long as *Cheap PC Bloke* enters in to an agreement stating that they will only sell PCs with *CAB*'s product. Furthermore, the agreement also states that should *Cheap PC Bloke* sell their product without *CAB*'s OS, they will be liable to pay 100% full market value for all *CAB* products.

Now *Cheap PC Bloke* really want this market edge and happily

agreement like this with Microsoft, but it wouldn't surprise me as I am sure there must be many more companies with an agreement that is similar to this than I have fingers and toes.

But don't blame Microsoft for these strong-arm sales tactics, it happens all the time all over the world with lots of companies in all areas of business, especially if they have the market cornered.

Trevor Henderson, *via email*

It would perhaps be a lot more honest if companies merely said it wasn't their policy to sell computer systems without an OS installed and leave it at that. I know it wouldn't be the first time someone had been misinformed by a salesperson (not necessarily maliciously in this case). It isn't an odd request after all. If I had a PC running Windows and it caught fire, I might want to buy a new PC, but I wouldn't want to buy another copy of Windows. We wouldn't buy new copies of any other software we already owned, so why should we have to buy an operating system with each new unit of hardware?

PseudoOpenSource

As a Linux user of several years, and an experienced software developer, I am very interested in writing Open Source code. However, I cringe at the thought of Windows users downloading and using it. I think software licensing should reflect OS licensing, *ie* if your OS is Open Source, use Open Source software – if it's closed and expensive, pay for your software.

This is where *pseudo open source* comes in. Is it conceivable that an OSS license could be

penned allowing developers to release a 'per OS' license, whereby proprietary code could be ported to Linux and released as Open Source while still being proprietary on Windows and other closed systems? Collaborators would have control over whether or not their code can be included on the closed system.

This would allow development firms to port their systems to Linux (*Dreamweaver* on Debian?) while keeping their existing revenue streams intact. In return, they would receive an unparalleled level of feedback and fixes from the Linux community.

I believe such a license would open a huge range of applications to the Linux user, and offer Windows users yet another cost-saving incentive to migrate.

Thanks to Borland Kylix porting has never been easier – provided you use *C++Builder* or *Delphi*. With *pseudo open source*, Borland's competitors would be encouraged to follow suit – so that even if developers kept their Linux products closed, they could still be easily ported.

I'd love to hear any feedback or advice. Keep up the good work Mark Healy, *via email*

Well, it's an interesting idea, but not one that would get the backing of groups like the Free Software Foundation, which is about giving users freedom. Obviously though, some companies do operate dual licensing policies (for example TrollTech with Qt). There isn't anything stopping you creating a licence of this kind. What do other developers and readers think of this idea? Let us know at the usual address. [LXF](mailto:lx@futurenet.co.uk)

SUBMISSION ADVICE

WHAT WE WANT:

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

WHAT WE DON'T WANT:

- Technical questions – direct those to our Q&A pages!
- Random abuse
- Nonsense rants
- 200 pages of meandering diatribe

WRITE TO US AT:

Linux Format, Future Publishing, 30 Monmouth Street, Bath BA1 2BW or email: lx@futurenet.co.uk

Reviews

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

LXF verdict explained

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

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

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

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

Value for money: Does it have a competitive price?

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



10 The close to perfect product.



8-9 Good, but has a few niggles.



6-7 Does the job, but needs work.



5-4 Average.



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

The Top Stuff Award

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



WHAT'S NEW...

SuSE Desktop

Office without Microsoft? Corporate users would be hard-pressed to tell the difference — except in the price! **p20**

Via Eden ITX >>

Size IS important — though it's a mini-ITX, it'll fit in a normal case if you need it to. There's even one of these wall-mounted in an interactive picture frame **p23**

Dantz backup >>


Though it won't run entirely on a Linux setup, you can still do backups of Linux systems on a network with this **p26**

Lycoris update

Not just a distro — these folks from Redmond are hoping to emulate their larger neighbour by providing a complete OS solution, including their own iteration of *OpenOffice.org* **p28**

Books

DNS & BIND Cookbook; Red Hat Linux 8 For Dummies; Red Hat Linux Bible; R Stallman's Free Software, Free Society **p31**



COMING UP SOON...

Kapital 1.1

The Kompany's personal finance software didn't arrive in time to make this issue

Gnucash

You wait ages for a personal finance app, then two turn up! Which one will be the winner?

eScape

Cheap desktops with pre-installed Lindows software hit the UK with Evesham's eScape series.

Opera 7

Well-publicised Windows-version bugs are holding this one up, let's hope Linux wrinkles are ironed out

Scribus

DTP with Linux is coming on leaps and bounds. Check out Scribus if you don't believe us!

JRun 4

The very latest version of Macromedia's popular drag-and-drop-capable J2EE server

LINUX DISTRIBUTION

SuSE Office Desktop

Hot on the heels of SuSE 8.1 comes this value-added distribution aimed at the corporate/home desktop. Nick Veitch puts it to work...

Desktop optimised distribution including commercial extras. Also consider Mandrake 9, Red Hat 8 or upgrading your current distro with Codeweavers software.

- **DEVELOPER** SuSE
- **WEB** www.suse.co.uk
- **PRICE** £89

The last issue of *Linux Format* contained a comprehensive feature on Linux's state of readiness for desktop use, and concluded that for a lot of people, it's useable now. This new themed version of SuSE Linux is aimed at increasing the size of that group, by targeting the average computer user who wants office applications and connectivity, but has little interest in the command line. The distribution is effectively a modified version of SuSE 8.1, with the removal of some superfluous software and the addition of some proprietary tools for common specific tasks.

Installation is similar to the previous SuSE distro (see LXF34) and should present no significant problems for those with a smattering of computer knowledge. The exception perhaps being where dual boot systems are required, but SuSE have provided for this by including Acronis

OS Selector (see the box opposite) that can cope with resizing partitions and making space for Linux.

The four CD set includes only two install CDs. In addition to the Acronis disk, the other slot in the case is taken up by the CD of source code.

Desktop

Once installed and rebooted, the wealth of Linux is open for business. This distro includes both KDE 3 and GNOME 2, giving a choice of user environments. Each has been customised with the addition of SuSE specific shortcuts, panels, menus and desktop items, so whichever you choose doesn't really make a difference from a functional point of view. Choices of software for specific tasks have been rationalised to present only a few options. This throws up a few oddities. Although *Mozilla* is included, it doesn't appear on any menu by default – *Konqueror* is the only web browser listed. This is also the case if you are running a GNOME desktop. At least this less cluttered makes things straightforward for the users though.

Crossover Office

Codeweavers have made lots of significant contributions towards *Wine*. *Wine* is a system that tries to replicate



Run Internet Explorer 5.5 and take advantage of Shockwave...

Included software

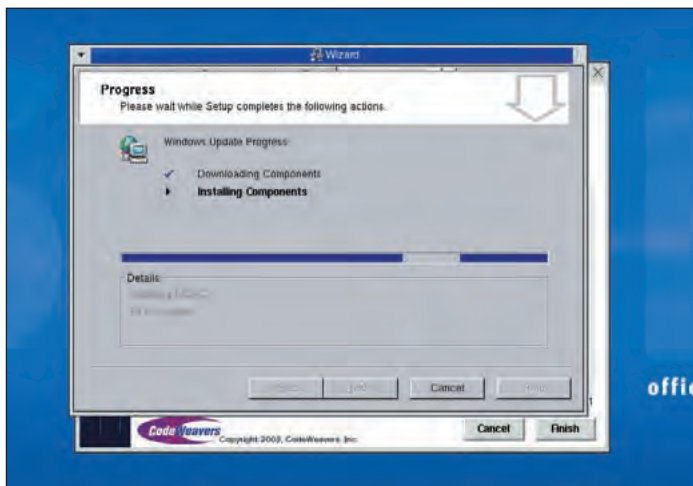
Kernel 2.4.19
KDE 3.0.4
Gnome 2.0.6
Glibc 2.2.5
Xfree86 4.2.0
Samba 2.2.5
Perl 5.8
Mozilla 1.0.1

the Microsoft API under X Windows, to allow Windows software to run on Linux. The project has been largely successful, but a lot of software, particularly software produced by Microsoft themselves, seems to stray out of the usual API framework and thus has difficulty running under *Wine*.

Codeweavers develop their own version of *Wine*, and also *Crossover Office*. This is a commercial product, which has been tweaked in order to get it to run popular Windows software, particularly the *Microsoft*

Office suite of software. If you have already bought *Microsoft Office*, you can simply install it on Linux, using the *Crossover Office* installer. All of the individual components will then run exactly as if they were running under Windows. Supported software under *Crossover Office* includes *Microsoft Word 97/2000*, *Microsoft Excel 97/2000* and *Microsoft PowerPoint 97/2000*. *Access*, the database software from the *MS Office Suite*, works to some extent, but is not officially supported.

Why would you want to run *Microsoft Office* if you already have *StarOffice*? One of the things currently not possible with *StarOffice* is the use of VBScript macros embedded in *MS Office* documents. Many *Excel* users deploy very complicated scripts in worksheets for a variety of tasks that would be impossible or at least time consuming to do another way. Rewriting these to work with *StarOffice*'s macro language may not be desirable. Plus, obviously, a lot of people are familiar with the



Installing Windows software, courtesy of Codeweavers Crossover.

Acronis OS Selector

Controlling your dual boot

SuSE have recognised that in the target market for this version of their Linux OS, many people may want to dual boot with a Windows system, at least until they get used to Linux. The disk partitioning software used for the SuSE install can't resize NTFS partitions though, so if you are using NT, 2000 or XP with NTFS partitions, you'll have to resort to other methods to make the space.

Bundled with SuSE Office Desktop is Acronis OS Selector – a boot and partition manager which will resize all sorts of partitions (including various

shades of FAT, NTFS and Linux partition types, including Reiser, Ext3 and swap partitions).

The OS Selector installs from a self-booting CD. You will need to already have a Windows installation or at least 10MB of unused disk space, as the software requires some FAT partition space to work from.

As well as being able to resize your Windows partitions, it can help manage other OS installs too, including multiple versions of Linux, Windows, BeOS, BSD, OS/2 and will even have a stab at OSes unknown to it.

likes of MS Word and would be happier to continue using it.

Other applications supported by Crossover Office include Visio, Quicken and Internet Explorer 5.5. With the included Crossover Plugin software, it's also possible to download and use a number of web plugin technologies such as Shockwave, Flash 6, Quicktime and even Windows Media Player. Other Windows applications could potentially be installed, but they are not actively supported by Codeweavers, so if you run into trouble with them, it'll be up to you to find your own help resources.

StarOffice 6

Many words have already flowed over the pages of *Linux Format* on the subject of StarOffice 6 and its free relative, OpenOffice.org. In our ultimate office roundup back in LXF30 this duo walked away with top honours thanks to the feature rich, easy-to-use nature of the software.

There is one slightly disappointing omission from the suite is a desktop-style database application (like MS Access). StarOffice 6 does include the adabas database server, but this is of the SQL/server style type of database, which probably doesn't equate to most desktop users expectations. It can still be used as a data source though, along with a number of other databases including DB2, ODBC compliant sources and even your Mozilla/Netscape address book.

The rest of the software is fantastic though. StarOffice 6 includes plenty of templates and useful support files not included with OpenOffice.org, so the differences are more than just cosmetic. One of the aims of this

version of SuSE is to create a desktop Linux that can easily be deployed in small businesses or the home. This means networking, and probably dealing with Windows servers or desktops as well.

Networking tools

Linux already has the tools for this. Samba provides more than adequate support both for running a Windows compatible shared folder on a Linux machine and connecting to a shared resource on a Windows box. Configuration isn't always as straightforward as on Windows though.

That's one of the reasons SuSE Office Desktop includes a special desktop resource to perform task such as install Windows software (which launches CrossoverOffice), share directories with other users (which uses NFS for some reason), setting up printers and connecting to other shares. Unfortunately, few of these are

as straightforward as they might seem, bar the Crossover Office option.

On close inspection, it seems that the Samba server is not included. Shares are set up using NFS. This is fine if you are sharing with other Linux users (security implications notwithstanding), but no Windows desktop will be able to see them. This isn't explained and could be the source of some confusion.

Connecting to Samba shares can end up being a bit of a lottery too. In our test the YAST tool only displayed a handful of the known shares available on the network, possibly because of its size and diverse 'Workgroups'. We were able to list, connect and use the others manually with the smbclient software, but not through YAST.

Printer setup is a little easier, though you'd be advised to use a CUPS server or direct IP printing than trial-and-error fiddling around with Windows printer queues.

Packages

One of the rather strange aspects of this release is the strange selection of

packages that are included. As it is designed as a desktop, it isn't surprising that server stuff such as Apache hasn't been included. Similarly, there's no MySQL, though there is a version of postfix included. Stretching the 'desktop system' angle is the omission of gcc. OK, so the average desktop user might not do much program development, but when the poor sysadmin comes to install some needed service or app, he's going to have a hard time without a compiler. Oddly there is a section in the software manager for C/C++ development. Taking a quick look reveals that some development tools are actually included in this release, such as the gdb debugger – a little odd to have the debugger without the means to compile any code!

Most of the omissions are sensible though – It's easy to bloat a system out including everything bar the GTKitchenSink which may only serve to confuse the intended user.

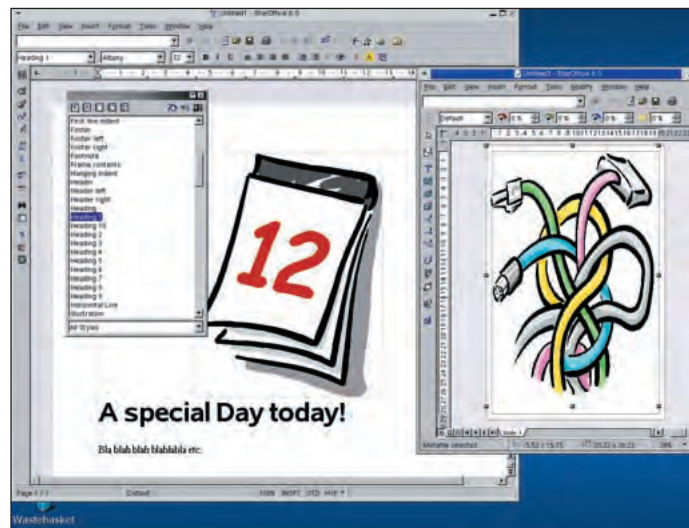
Great Value

At £50 more than the Personal Edition of SuSE 8.1, you do actually still get value for money with SuSE Office Desktop. StarOffice as a standalone product costs around £35 if you shop around. And then you have to go through the hassle of actually installing it. The Acronis software retails at around £30 and the two Codeweavers products would set you back US\$80 if ordered online. Including all these for under £100 offers good value for money.

It would be nice to conclude that this is the ideal solution for office/desktop use, but some of the omissions and difficulties cannot be ignored. None of them are insoluble though. A decent tool for setting up Samba sharing, a look at some of the YAST modules and a slight rethink on the standard packages included and this could be a winner. **LXF**

Documentation

The Office Desktop is accompanied by its own manual covering a variety of topics. As well as installation and basic configuration, there are sections on the finer aspects of StarOffice, working with Crossover Office and some of the more prominent Linux tools included. With over 350 pages it does deliver a lot of useful information for the Linux beginner or first-time SuSE user.



StarOffice - the best office suite available for Linux today.

VERDICT

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

Overall a great effort, with just some niggling problems holding it back.

LINUX FORMAT RATING
8/10

MINI ITX MOTHERBOARD

VIA EPIA M9000

A fully specced motherboard the size of a software box? **Nick Veitch** examines the diminutive new board from VIA

miniITX motherboard. Slower processor version available. Also consider ESP 5000.

■ **MANUFACTURERS** VIA Technologies
 ■ **WEB** www.viapsd.com
 ■ **PRICE** £125

The last time we reviewed a mini-ITX motherboard we got a lot of feedback from people who had them, or just wanted them. The tiny form factor belies the features of these boards. They may not offer much in the way of expansion – but what would you add to a board that has it all?

Connectivity

The VIA team hasn't skimped on connectivity – important if the box is to be taken seriously as a media hub device. For onboard devices there is a single floppy controller and 2 UltraDMA 133/100/66 connectors. The usual collection of IR, I2C and audio headers litter the board, nestled in amongst the chips and other components. There are so many connections packed in so small a space that the circuit router must have been working overtime.

For external devices, the board sports two USB 2.0 connections on the rear panel and a further two headers inside, which can also be routed to the real plane with the backplane interface provided (though depending on your case, this may

obscure the PCI slot). USB2.0 is not very rigorously supported in current release kernels, but it will be in 2.6. Of course, the ports are backward-compatible with USB 1.1 devices, so you can still make use of those safely.

In addition to USB2, the board also has two IEEE1394 ports. This could be crucial to many media hub applications. Although many peripheral manufacturers have opted for USB/USB 2.0 for the likes of scanners, printers etc, the world of Digital Video is still firmly entrenched in the i.Link/Firewire/IEEE1394 camp. The controller is the VIA VT6306 chipset, which is fully OHCI compliant and will work happily on any IEEE1394 enabled kernel.

A LAN connector is onboard too. This is part of the VT8235 chipset and Via produce their own driver for it, supplied as either a pre-compiled module or as source for integration into custom kernels. The standard kernel will detect it as a Rhine chipset device and use it quite happily though. Importantly this has thin client support so you could use it as the basis of a diskless booting system connected to a server.

On the memory side of things, a single DDR DIMM slot will easily accommodate a single 64,128,256,512 or 1GB module. This may seem a little excessive – it probably isn't going to need the performance of DDR for most uses, and older SDRAM DIMMS are slightly cheaper still, but this is mainly because of the chipset support for

DDR. As it is pretty much a standard on new boards, this doesn't seem too much of an imposition.

Sound

The Sound is based around the VT1616 six channel audio device – to fit in with the 'multimedia' theme it has onboard 5.1 sound. The back panel spec only allows for three audio jacks, which are usually designated as line-in, »



We tested several distros on the VIA EPIA mini-ITX board, but defaulted to Red Hat 7.3 due to driver issues.

Reviews VIA ITX Motherboard

line-out and Mic. VIA has had the cunning idea of dual-purposing these connectors as jacks for centre, front and rear speakers for users that like to take advantage of surround sound provided by DVD movies and some games, though this is not supported or enabled on Linux yet.

VT1616 support has been added to the 2.5.x series kernel, so it's likely that it will be fully supported in 2.6. However, as it stands the audio will work perfectly well in 'normal' stereo mode with the usual AC97 drivers.

Video

This board features onboard accelerated 4x AGP graphics, with an included hardware MPEG2 decoder. Getting it to work might prove a bit of a struggle though. The drivers required are not part of the standard XFree86 distribution, so if you are doing a normal install, you'll only be able to select standard VESA graphics modes. It works fine as it is, but unfortunately you won't be able to enjoy the full potential of the hardware – yet.

VIA does provide its own drivers and libraries to use with the CLE266 hardware though. The only drawback is that the ones available are precompiled for specific versions of Linux, and these haven't been created for the latest range of desktop distros. Mandrake 8.2 and various shades of Red Hat 7.x are supported. And if you think you can force the module to load on a different kernel, don't bother – these were compiled with *gcc2*, and both Mandrake and Red Hat now use *gcc3* compiled kernels. We had to revert to Red Hat 7.3 to test out the drivers for this review.

However, if you are using one of the supported distributions, installation is pretty painless.

The onboard graphics chipset also delivers a TV-out signal – essential for home media and many embedded applications. This is switched as either S-Video or RCA (composite) output, so it's suitable to be connected directly to most video devices and TVs. When we reviewed the Eden-based ESP5000, setting up an X display mode that would work with the TV-out proved almost impossible. This board uses the new VT1622 TV-out device which is much more amenable. It supports PAL or NTSC modes at a variety of resolutions including 720x576, with support for under and overscanning the image to get the best results. VIA provide detailed instructions on X configuration settings on the www.viaarena.com site, along with a host of other useful information for Linux users.

There's also an LVDS header on the board if you want to drive some other type of display for embedded applications.

Processor

The board we tested featured a C3 processor running at 933MHz. This is really more than enough processing power for most desktops, and even media hubs. You probably won't be playing *Unreal Tournament 2003* on it – well, not without turning off a lot of the graphics options anyway. Linux reports 1861 Bogomips for the processor, which isn't too shabby.

There is also provisional support for VIA's 'Longhorn' features of this chip. This is a power-saving mode which

What on earth is Mini-ITX?

Good things come in small packages...

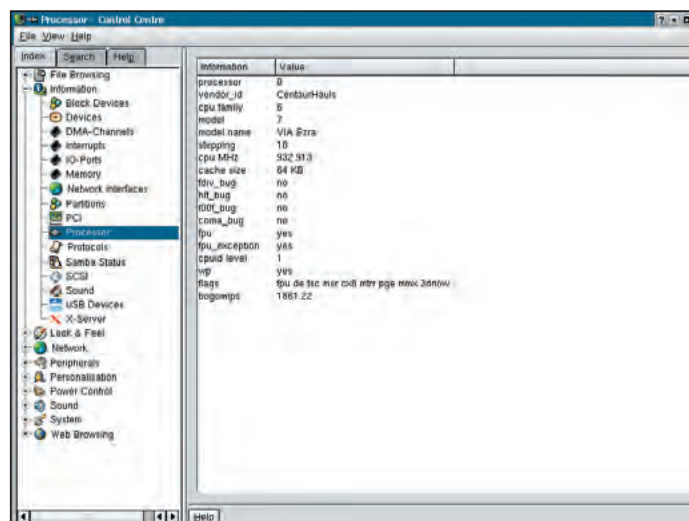
Mini-ITX is a mainboard form factor that follows on from such designs as Flex-ATX and Micro-ATX. The formfactor defines a number of properties for compliant boards, most importantly in this case, the physical size of the board.

Mini-ITX boards are 170cm square. That's pretty small. In fact, it's barely larger than the backplane with all the connectors on it. Another spec is that complying boards draw less than 100W of power. That's not a lot.

The other bonus of this spec is that you don't actually need specialist cases or PSUs to use it. The board has been cunningly designed to use some of the same mounting holes as the ATX spec. The backplane is the same, as is the power connector. The upshot of all this

is that the boards can be installed in a bog standard case if you like.

All of this leads one to the conclusion that these boards are designed for some specific purpose, but actually there's all sorts of things they can be used for. Cheap diskless clients, always on home servers, media hubs and a whole range of embedded applications. The low power consumption means less heat, making case designs easier, and in the case of Eden based boards, fanless too. You can discover a whole range of projects people have used miniITX boards for at the community site: www.mini-itx.com Robots, cars, and even a computerised picture frame are among some of the projects listed (and it's a good place to buy a selection of mini-ITX related goodies too.

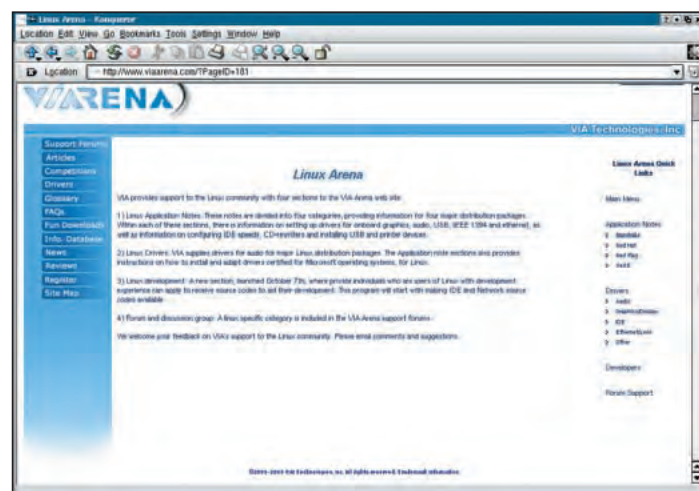


The processor is pretty much fast enough for most purposes.

lowers the frequency of the processor, to save power and reduce heat output. This is still pretty experimental, but the kernel module code and tools are available if you want to try it out. The on-board sensors, used for monitoring temperature and the like, also have provisional support in the *lm_sensors* kernel module.

Conclusion

This is a tremendously powerful little board, more than up to the task of being a 'digital-hub' device, and with plenty of extras useful for many embedded applications. The lack of up to date driver support is a concern. Some drivers for Mandrake 9 appeared during the course of review, so it seems VIA are now on the case.



There's plenty of info and drivers to download at the VIA Arena site.

VERDICT

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

More up-to-date driver support would be handy, especially to take advantage of the audio and graphics capabilities. Otherwise a fantastic board.

LINUX FORMAT RATING
9/10

BACKUP SOFTWARE

Dantz Retrospect Professional 6

Nick Veitch assesses this easy-to-use cross-platform backup solution

Straightforward backup to a range of hardware devices. Also consider BRU, Arkeia, Netvault and others.

- **DEVELOPER** Dantz Software
- **WEB** www.dantz.com
- **PRICE** from £88

Although the *Retrospect* backup server doesn't run under Linux, it can still be used to backup your Linux servers or workstations in one of two ways. The first is using the client software. Linux clients are included for networked backups. Simply install and run the client on the target system (it is supplied as an RPM and as a tar archive) and the backup server will be able to communicate with it to schedule backups and transfer data across the network. The tar file version of the client software comes with a simple installer, but it seemed to work on all the systems we tried (the software is x86 only for Linux however). A daemon running on the client machine then listens out for instruction from the backup server, or can signal a request for an immediate backup of data. A Java client app on the remote client can be used to monitor the status of backup activity, and signal a request for an unscheduled backup.

The use of clients on Unix systems requires a multi-server licence for Retrospect, but the other way of creating backups is through the use of shared folders over the network. Linux filesystems can be easily set up as Windows share using Samba. As long as the server can connect to them, they can be backed up.

Devices

Dantz *Retrospect* 6 supports a wide range of backup devices, including tape streamers, autoloaders, file backups, CD-R, DVD-R and DVD+R. Previous versions of the software received

criticism for the narrow range of drive supported, particularly for backup to CD or DVD, but the list has grown considerably since then, and includes a healthy range of SCSI and ATAPI devices, as well as IEEE and USB equipment.

Of course, as the backup hardware is supported only on the Windows platform that the server runs on, you need to be sure you have relevant drivers for that system. The backup server automatically recognises hardware on the system, so it isn't necessary to manually configure specific devices (if you only have one CD-RW drive and you select to backup to CD, that's the one that gets used).

In use

For straightforward backups, there is no need to configure a lot of options in advance. It's possible to run the server and just click on 'Backup'. The user will

then be taken through the aspects of configuring the source, defining a backup set and what selection filters will be used before proceeding. The selection filters can be user defined and include various file pattern matching and attributes. For Linux and other Unix OSes, this includes the ability to filter on permissions, and handling of symbolic links. You can create and manage a list of different filter sets, which can be as simple or as complex as you need.

For offsite storage, and when backup sets may be kept in an insecure location, encryption of the data is a good idea. *Retrospect* 6 supports two methods of encryption and a simple password mechanism to protect data. The most

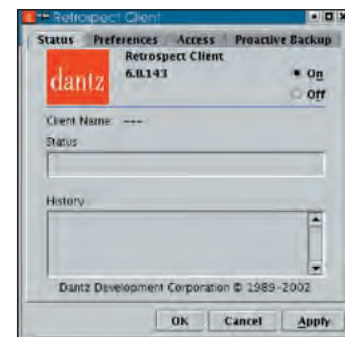
elaborate encryption is DES which, while not as robust as encryption methods supported by Linux, is better than none.

For more advanced backup regimes, there is full support for backup set creation, including multiples sets for rotation. In fact, there is an *Easyscript* wizard to create a complete schedule with rotation of all backup sets. Backups default to using an incremental strategy, saving only the changes since the last full backup, but this can be adjusted to repeat full backups on a regular basis if that's what is needed.

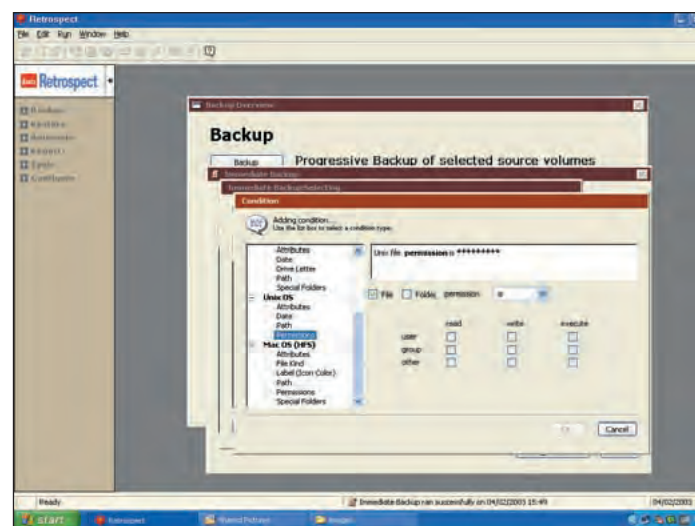
Restore operations include single file restores and a disaster recovery mode – useful if used in conjunction with removable media backups.

Conclusion

Dantz Retrospect combines an approachable ease of use for those with simple requirements, with a feature laden capacity for more complex and



The Java client app can be used to set local preferences and request an unscheduled backup.



Fairly complex filters can be used to back up only the data required.

precise backup management. The lack of a Linux server application means it will only be suitable for those running heterogenous network environments, which is a disappointment. **LXF**

VERDICT

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

Good value backup solution for mixed environments, which is also easy to use. Better encryption would be preferable.

LINUX FORMAT RATING
8/10

OS, DESKTOP AND OFFICE UTILITIES

Lycoris ProductivityPak

In their quest to bring a complete productivity system to the non-Microsoft desktop, Lycoris have released an office suite based around *OpenOffice.org* and integrated it into their Desktop/LX operating system. In search of OO heaven, **Andy Channelle** takes it for a spin...

Office suite which faces stiff competition from the likes of *KOffice*, *OpenOffice.org* and *StarOffice*.

DEVELOPER

Lycoris/OpenOffice.org

PRICE

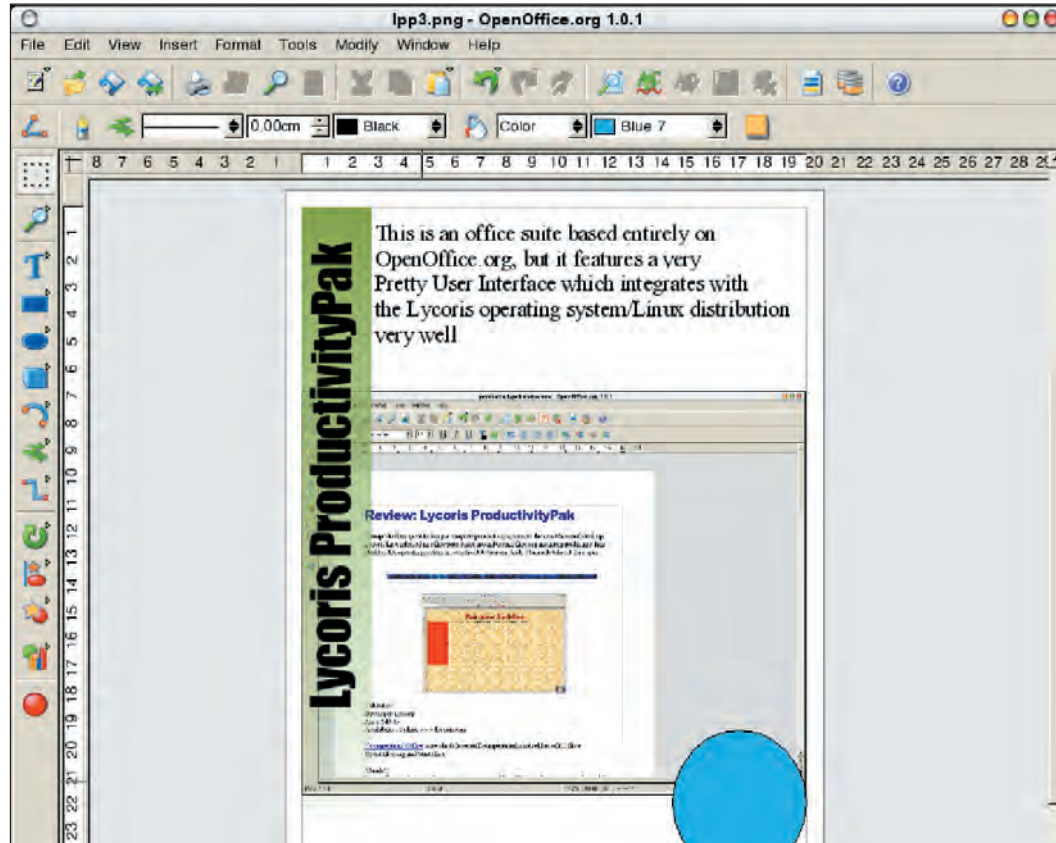
\$49.95

AVAILABILITY

Online,
www.lycoris.com

Lycoris has ambitions to be more than just a run-of-the-mill Linux distribution vendor and it doesn't appear content to rely on 'services' as a revenue stream. In fact, the company - who used to be called Redmond Linux (hint, hint) - hope to emulate another Operating System builder in their area. To that end, Lycoris don't sell a Linux distro, they sell an operating system designed to cover the essentials for the home or small office. These are second-generation resellers and their strategy is to improve the usability factor by trimming the fat or, as some would claim, removing choice.

Productivity-wise, Desktop/LX includes *KOffice* as its core application set, but Lycoris also now offer an upgrade in the form of their own tweaked version of *OpenOffice.org* which they've renamed the *Lycoris ProductivityPak* (*LPP*) and slapped on a US\$49 price tag. But what can



Lycoris Artist covers all the visual elements of your productions.

ProductivityPak offer that its free counterpart can't? And is it enough to convince the average user to part with their money? The first step on the *ProductivityPak* upgrade is ensuring

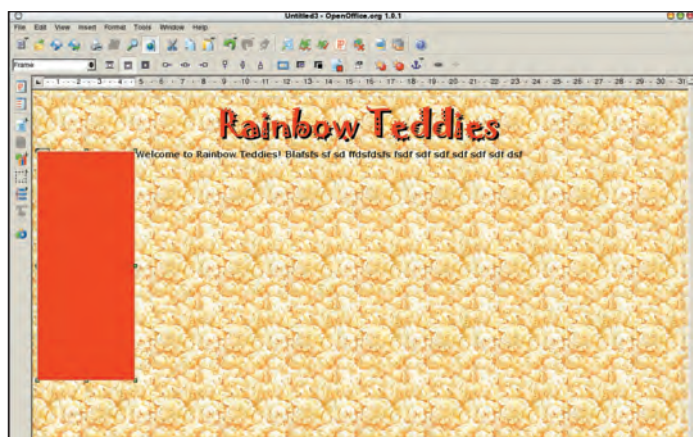
you're running the most recent (Amethyst) version of the Desktop/LX operating system. As this merely involves launching the Update Wizard - for registered users at least - it is not a problem.

Integration station

Consulting the single page of documentation reveals that next you insert the CD, which automounts and then opens *Konqueror* to display its contents, and then tells you to click on the *ProductivityPak* icon (which is different than in the leaflet, which might bamboozle the odd novice or two!) This kicks off the installer and within a couple of minutes you have a fully useable office suite.

The installer is the usual slick *OpenOffice* effort with a novel twist: once the suite has been installed, the script

discards the *KOffice* entries from the Flower/Productivity menu and updates them with *Lycoris Artist*, *Lycoris Writer* etc. A cursory glance through *KPackage* suggests the *KOffice* suite hasn't been completely removed, but to all intents and purposes it is not there. However, if you double-click on, for instance, a file with the .txt extension it will open in *KWord*, which seems to be a lost opportunity; *Lycoris Writer* doesn't even appear as an option in the 'open with' flyout. If your installation script can handle big changes on the *K Panel*, it should be able to update the file associations with relative ease. All the main Microsoft file formats default to *LPP*, so perhaps Lycoris opted to leave the more mundane extensions to more basic and, it must be said, faster loading applications.



The Lycoris word processor is as much at home with web pages...

As *LPP* is designed primarily for created printed work, I decided to do the integrated printer setup before actually getting to grips with the applications. Disappointingly there was no option for my aged Epson StylusColour 400, so I chose the nearest model (an 800) which gave adequate results. Opting for the generic printer brings up the usual *KPrint* menu and proved to give better results in a quicker time.

The suite is divided into four applications – *Lycoris Artist, Presenter, Spreadsheet, Writer/WebWriter* – and a pair of setup applets covering printer configuration and suite removal/repair. So far, so *OpenOffice*.

Launching any of the applications reveals the biggest change Lycoris have wrought on *OpenOffice*. Instead of the standard *OOo* look and feel, which was very grey and *Office97*, Lycoris has grafted a new set of icons onto the user interface which nicely follow the conventions used in the OS. It makes for a very integrated experience, but the icons do look a little reminiscent of Windows XP. Though some users may find the look a little twee, I like it and as it doesn't obscure any of the options the added colour does make the screen a nicer thing to look at. The redesign is admirably consistent with, for some reason, only the 'New Document' icon and help system left unreconstructed.

Applications

LPP covers all the tasks that the average office and home user would need, and it does them all well and within a very consistent environment.

The heart of any office suite is the word processor and Lycoris has given two slots on the *KPanel* to this task. *Writer* launches the WP with the standard tools for printed output, while *Web Writer* opens the same application configured for HTML work. It's not a big thing, but is a nice touch.

The word processor itself is powerful, easy to use and is built on the idea that the majority of users will have basic requirements, which are almost all met on the 'first layer' of the UI. Power users will find that the more specialised options are still just a few clicks away, but concentrating on the core 20 per cent of features that get used 80 per cent of the time has resulted in a clean and well balanced UI that doesn't dominate too much of the screen like some other WPs.

As the Desktop/LX installation uses Freetype2, and had linked to the TrueType fonts already installed on hda1, access to typefaces from within the suite – often a cause of cross-platform problems – is trivial and really eases the transition from Windows to Linux. All of the UI elements feature anti-aliased text as well which makes it more readable.

The only real disappointment here – and the one thing Lycoris could do to put a little distance between itself and the standard *OpenOffice* download – is the lack of localised dictionaries. You can, of course pay a visit to www.openoffice.org and download a UK (or whatever) dictionary, but it would have made sense, and not taken too much work, to throw in a few of the more common dictionaries and give the product a unique selling point over its no-cost progenitor. As it is, you'll have to download and install your own dictionaries.

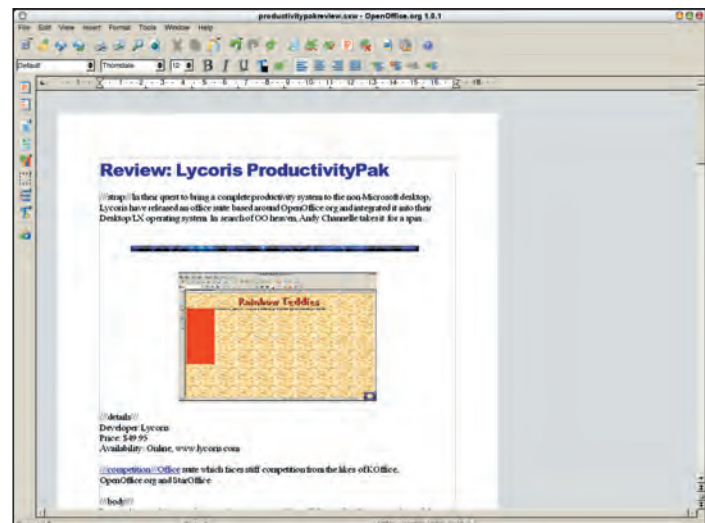
The rest of the suite does everything the *OpenOffice.org* product does, and as *OOo* is the winner of many awards and head-to-heads, you can be sure you're getting a set of applications that will meet or exceed your expectations. It's as much at home organising your accounts as it is preparing presentations or building a website. Indeed for much of this work you may never need to leave the suite, which is a pretty big endorsement.

The integration between the applications is seamless. For instance, you can cut and paste an image from the *Artist* module, paste it within a WP document and then edit in place. The menubar and top toolbar remain the same, but a subset of the drawing tools appear on the second level toolbar. Having the appropriate tools to hand at all times makes for an efficient working environment and would encourage users to experiment with style and layout.

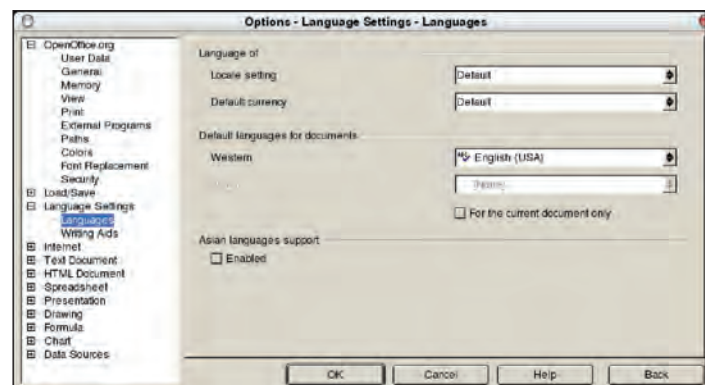
One thing I liked a lot is the fact that once the *ProductivityPak* was installed, the Lycoris Update Wizard appeared to look for additions to the suite as well as the core operating system, which should make updating to the recently released *OOo* 1.0.2 a joy rather than a chore.

Conclusion

Although altering the look of the application is an entirely superficial change, it fits in nicely with the Lycoris ideal of giving the user a consistent



...As it is with its traditional paper-oriented tasks.



Some other language options beyond American English would be nice.

interface, whether they're browsing the web or creating a presentation. The other big change is in the area of support. Although Lycoris will never offer the kind of options that Sun can provide for *StarOffice*, there is at least something available beyond the help files. For the money though, a decent printed manual containing, perhaps, a range of tutorials on common tasks would have been a bonus. As it is we're stuck with the unadulterated help files, which are fine, but again this would have been a great opportunity for Lycoris to give some added value and make the product more appealing to the less technical buyer.

As for the applications: they simply can't be faulted built, as they are, on the comprehensive *OpenOffice.org* project. The more complex MS files may need a little work when imported but unless you have thousands of such documents *OOo* is more than capable of being a day-to-day workhorse. The only issue appears to be the cost which, compared to *OOo*, may be a factor. But if you consider that the you could equip four PCs with Desktop/LX and the

ProductivityPak for the cost of one Office XP licence, it looks a bargain.

Lycoris have crafted a very coherent operating system and office suite combination that is suitable for general home use or in an office environment. The look is very contemporary and the feature-set would appear to cover almost every need. A few more tweaks – especially in the area of documentation – would make *ProductivityPak* a more attractive option for corporate buyers or those equipping a small office. **LXF**

VERDICT

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

Lycoris has its work cut out competing with *OOo*, but this is a slick, solid and well integrated solution that should appeal to both home and office users.

LINUX FORMAT RATING
8/10

DNS & BIND Cookbook

Chris Denton goes through a handy reference collection of tasty nameserver recipes.

■ **PUBLISHER** O'Reilly
 ■ **AUTHOR** Cricket Liu
 ■ **ISBN** 0-596-00410-9
 ■ **PRICE** £24.95

O'Reilly's weighty *DNS & BIND*, 4th Edition, is the definitive reference guide for BIND administrators; but often its very authoritativeness makes it difficult to simply dip in and find out how to perform various commonly required tasks. Too often, you end up wading through an entire chapter or more just to find the one setting that solves the problem. The *DNS & BIND Cookbook* is new companion volume offering a shortcut. It is presented as a series of objectives, divided into themed sections, with brief descriptions of the steps involved in achieving them.

Beginning with useful advice on implementing BIND in the first place, the recipes get progressively more

sophisticated. There are chapters dealing with zone data, email, IPv6, logging and troubleshooting to name a few. Particularly impressive is the coverage of security. Like any Internet program, BIND requires close and careful attention to ensure that vulnerabilities are not inadvertently left unattended. *Cookbook* details many key areas of concern, such as concealing the version number, working with a Firewall and protecting against spoofing. Any sysadmin that has to make BIND co-operate with Windows DNS will relish the recipes dealing with interoperability, as will those concerned with Microsoft's implementation of dynamic DNS. In fact, so many areas of nameserver admin are covered there will almost certainly be something of value to anyone running BIND, no matter what the setup.

If this book has a weakness, it's the brevity of each recipe. The majority of them are a couple of pages long at best, and this often does not allow enough



space for a proper walkthrough of solutions. Instead, just enough info to get by is offered, with further research required, either in other recipes or the main *DNS & BIND* reference book. Only after a bit of index-trawling are some of the more complex recipes in any way clear. While this can be annoying, the need for countless repetitions of the same info is thus avoided and at no time does the shortness of a particular solution make it impossible to utilise it. Although by its nature fragmented, this book should not be thought of as a loose collection of HOWTO documents.

The author, Cricket Liu, is also co-writer of *DNS & BIND*, 4th Edition, so the depth of his knowledge is

unquestionable. *Cookbook* keeps well away from DNS theory, and is all the better for sticking closely to the practicalities of BIND. A valuable resource for old hands and BIND newbies alike, anyone running the world's most widely used nameserver is likely to find the *DNS & BIND Cookbook* nigh-on indispensable.

VERDICT

Not quite everything there is to know about running BIND, but pretty close.

LINUX FORMAT RATING
 9/10

Red Hat Linux 8 For Dummies

Tom Wilkinson investigates whether or not Red Hat Linux 8 really is for dummies...


■ **AUTHORS** Jon "maddog" Hall and Paul G. Sery
 ■ **ISBN** 07645-1681-7
 ■ **PRICE** £22.49

This book aims to educate those who have heard about Linux, but don't know much about it. Starting with an intro of what Linux is useful for, the book sets out in a clear and informative manner why you might want to run it in place of or alongside Microsoft Windows.

Starting with hard disk repartitioning to allow space for the install, the book guides new users through defragging the disk and using the FIPS partition splitter. This is followed by a detailed step-through of all the questions asked during install. A small section is devoted

to choosing a good password even at this early stage, something that not enough users consider important. Between sections on using the system and installing, some differences between Windows and Linux are introduced – the filesystem tree and users.

Setting up Internet access comes before sections describing apps that come with the system – once Internet access is set up, it's easier to download help, updates or add-ons. A tour of the GNOME desktop – or rather BlueCurve, Red Hat's customised version is next. Instructions on the use of various applications, including productivity and multimedia are included. These include links to sites which allow playback of MP3s and Ogg Vorbis under the system – something Red Hat removed due to patent concerns, allowing users to copy



their CDs to their computer.

About a fifth of the book is an intro to more complex aspects of Linux – the shell and package management. Where possible, graphical programs are used to demonstrate actions as opposed to the text interfaces – as this is a book aimed at novices, this isn't a criticism. Final chapters look into problem-solving, online help and man pages.

The book comes with the two-disc "Publisher's Edition" of Red Hat Linux. We're unsure as to what packages are missing compared to the GPL edition

which has three discs, however it does come supplied with all the free applications mentioned in the book.

VERDICT

A clear introduction to Red Hat Linux for the newcomer, with a friendly and informative writing style.

LINUX FORMAT RATING
 8/10

Red Hat Linux 8 Bible

Marco Fioretti introduces a small encyclopedia beneficial to many users and sysadmins.

■ **AUTHOR** Christopher Negus
 ■ **ISBN** 0-7645-4968-5
 ■ **WEB** www.wiley.com
 ■ **PRICE** £37.50 incl. VAT

This massive book (1.5 kilograms!) is composed of 1062 pages and three CDROMs, containing the whole Red Hat Linux 8.0 distribution. Its purpose is to teach to novice and less-than-novice Linux users how to do almost everything with Red Hat OS, from video conference to running a web server. The reader is not required to be a programmer (the suggested experience level is "beginner to advanced"). No complex tool should be used without some patience, good will, and experimentation; and this is where this *Bible* helps. The approach is task-oriented, with plenty of tables summarising commands, locations of useful files, and so on. This happens at every level, starting from the chapters titles: they are not technically named, not for example: *iptables internals* or *network daemons*; but much more accessibly, for instance *Getting your desktop to work*, *Monitoring system performance*, or *Administering a mailing list*.

Structure

The material is divided in twenty-five chapters grouped in four parts. The first one is predictably *Getting Started*, and contains two chapters explaining what Linux and Free Software in general are all about.

The second part moves to actually *Using Red Hat Linux* as a desktop system. One thing that I like here is that, even if the book has not the purpose to create hardcore command line gurus, it does start from there, and no previous experience is necessary. The audience is gently introduced to the concepts of command line interfaces: environment variables, aliases, scripts, common setup files, *vi* and many others.

It's important that readers who have always used only GUI based operating systems and tools should be

sure to have understood this first part before going on. Linux (and any U*nix-like system) can be certainly used in GUI-only mode today, but scripts and text-based configuration still rule under the hood. Ignoring them in your daily computing, and modifying them only through point-and-click interfaces is fine, as long as you know where to look when trouble comes.

After this console tutorial, a tour of the desktop environments available under Red Hat 8.0 begins, introducing KDE, GNOME, and the Red Hat way to configure them with the Bluecurve GUI, whose rationale is explained here <http://people.redhat.com/otaylor/rh-desktop.html>.

Running applications

The next chapter describes how to run both native and foreign applications (with emulators) under Red Hat Linux. The explanation of *Wine* (WINDows Emulator or Wine Is Not an Emulator) is clear and detailed. A nice tip that I was pleased to find in this chapter is the recommendation that users should bind any program to icons or menu entries only when they are sure that it's working without errors. Before that point, the program should be started by typing its name and options inside a terminal, which will display any error, and make it much easier to report them when asking for help! The second half of part two covers Publishing, Games, Multimedia, and the Internet, reserving one whole chapter for each of these topics. Publishing is succinctly described in all its Unix flavours, from the historical *groff* to the *dernier-cri*, *OpenOffice.org*, without neglecting the mighty *LaTeX*.

The review of the current state of games under Linux, be they native, or emulated Windows ones, includes a lot of pointers of online resources which give one less reason to reboot to 'that other' operating system. The same holds true for the *Multimedia* chapter, which explains how to watch TV, listen to music, manage digital photo archives, or set up video conferences with NetMeeting users. Part two ends with the *Internet And*



The Web chapter, which explains not only web browsing, email, newsgroup and instant messaging, but also illustrates how to connect securely to remote computers to run programs and copy files.

These first two sections (and, in general, the whole book) explain in detail only the program which the default Red Hat 8.0 setup allocates as first choice for that task – the book explains extensively only the first one or two entries in each default menu. This is for obvious reasons of clarity – there is so much Free Software available that to explain it all would take even more paper than the book already contains. Secondly, many readers will be expecting to be able to start *doing* things as soon as possible, and fiddle with choices and settings later, a bit at a time, as their experience and confidence in using the Red Hat Linux OS grows.

Administration

Part three concerns administration – how to configure and manage hardware, software, kernel modules and users. Everything is explained clearly, with the right level of detail, and plenty of examples and tables. The automation of system tasks, be they startup or periodical ones, together with backups and the use of Red Hat-specific administration tools are also covered. Chapter 14 is a good primer on the broad field of computer security issues.

The fourth and last section, *Red Hat Linux Network And Server Setup*, accounts by itself for almost half of the pages. Readers shouldn't fear

that the book is for IT professionals; after

all, Linux is more powerful than Windows, truly networked, multiuser and multitasking since (almost) since its first day. If all you wanted was an expensive game console with some word processing and web browsing capabilities you wouldn't be reading this review here, now would you?

More seriously, while beginners who may wish to learn system administration for a living will use every page of this section, a lot of it will come handy even to build a private family network to share files, printers, scanners, or an Internet connection. The first two chapters cover setting up a LAN (even a wireless one), and connecting it to the Internet in a secure way. Following this, there are eight chapters dedicated to the more important network services, from printing to DNS.

The book ends with the GNU GPL and three appendices. The first describes the content of the three attached CDROMs, the second lists all the RPMs packages included in (or removed from) Red Hat Linux 8.0, and the last one contains a series of tables summarising the use and setup of many network services.

The Pros

If you're a beginner who want to save some considerable time reading FAQs mailing list archives, the book is ideal. Reading the relevant chapter *before* asking for support online will also both help to understand *what* was

going wrong, and by sounding more informed, make your posted requests clearer, and therefore much more likely to get a timely response.

There are several other bits which I appreciated. To begin with, a clear explanation of the difference among the terms 'hacker', 'cracker' and 'script kiddie' is given: considering DMCA and friends, things like these are never explained often enough. Next, there is a nice bunch of white sheets at the end, so you can write down every setting (not your passwords!) Here are some other highlights:

- Appendix C, a nice network servers cheat sheet: a clear table lists all services, the related Red Hat package, its startup scripts, the path to the daemon binary, and all the related configuration files,

- The *iptables* tutorial, perhaps the best I know in its category (ie of that length, and for that target audience)

- The sections on the init procedure and the innards of cron, both explained very thoroughly and clearly

- The incremental backups explanation (maybe a bit too polarized towards tape-based backups, but many concepts are still valid when using CDRoms or external hard disks)

The Cons

On the negative side, I have to report some scattered glitches, mostly harmless. There are no colours, so the contrast on some snapshots is a bit poor. PCMCIA management would have benefited from a bit more space, and the index could be more complete.. (Of course, a really complete one would make another book all by itself!)

Two slightly more substantial issues I found are man pages and version discrepancies. In the first case, it is explained how to compose man pages with *groff*, but not how to read them, ie what the standard section and syntax mean. The novice reader would probably need that first.

Last but not least, there are occasional paragraphs that don't appear to have been updated from previous editions of the Bible. Here are some examples:

- On page 244 it reads: "other [graphic] tools that come with Red Hat Linux include xv". The next page says (correctly): "Because xv is shareware... is no longer part of the Linux distribution".

- On page 514: "Red Hat Linux sets up a firewall for you during installation using ipchains, so you have to disable it and enable iptables to use iptables instead" (true in Red Hat 7.x). In contrast, the very next page states: "in Red Hat Linux 8, *iptables* replaces *ipchains* as the firewall created during installation. To use *ipchains*, you must turn off *iptables*".

Please note: the actual explanations following these sorts of contradictions are always correct – not a problem for more experienced users, but Linux novices should take care to double-check if they are unsure.

Conclusion

In spite of the Cons, the *Red Hat Linux 8 Bible* can be a very good Linux starting point, especially to learn fun things and sysadmin responsibilities at the same time. It does a very good job

of mediating between the initial discomfort common to so many Linux newcomers, and getting them to learn and understand, possibly without pain, what one *has* to know, and then some.

When it comes to desktop applications, I personally prefer more spartan tools than the ones available by default in Red Hat's KDE or GNOME. Maybe this is why I will use part three much more than the others. Let me recommend one thing to all readers of this book: please do read *all* of the third part (and those chapters in the fourth related to services that you actually use)! Nobody will nag you if you don't surf with *Mozilla* at the greatest possible speed, but a thorough reading of those sections will save you a lot of time and, probably, grief.

VERDICT

All experience levels are catered for by this comprehensive and well-balanced mix of reference and instruction.

LINUX FORMAT RATING
 **8/10**

Free Software Free Society

Selected essays of Richard M. Stallman

Richard Smedley reads up on the ideals and purpose of the Free Software movement.

- **WEB** www.stallman.org
- **PUBLISHER** Free Software Foundation
- **ISBN** 1-8821-1498-1
- **PRICE** US\$24.95

Twenty years ago Richard M. Stallman began a personal quest. He was not prepared to live with software that could not be shared and improved, so set out to create a free operating system.

From the GNU manifesto, laying out the aims of his project, to a considered look back on the inexorable rise of Free Software, this is an assemblage of 20 years of careful, impassioned arguments on behalf of

freedom for all computer users. Heartfelt pleas against the folly of proprietary code are delivered in the form of humorous analogies and well-selected examples.

The collection does suffer from a seemingly hasty assemblage (the lack of index *may* be acceptable, but the poorly proof-read Editor's note is not); however Stallman's thoughts on copyright, community and coding are as clear and strong as ever.

These are not only notes on the social history of the world's hacking community and the development of not just our favourite operating system, but also a whole philosophy of sharing and cooperation. Here we find Stallman the concerned humanist, as much as Stallman the strident

denouncer of proprietary software, and of the siren Open Source which threatened to derail talk about freedom.

New writings on patents and copyright, and on current threats to freedom, add to the fairly well-known works collected, along with transcripts of a couple of speeches, and round off a first-hand account of the first 20 years of copylefted software.

Those who are still confused about the differences that exist between Free and Open Source, or why all licences are not the same will find a great deal of useful information here, and those


who do know will

download a copy to send essays to their friends, and keep a copy on their shelf for reference. **LXF**

VERDICT

A mixed collection of valuable original material from the founder of the Free Software movement.

LINUX FORMAT RATING
 **8/10**



Roundup

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

Audio Editors

OUR SELECTION AT A GLANCE

- Audacity
- Ecawave (with ecasound)
- GLAME
- GNoise
- mhWaveEdit
- Sweep

The continuing craze for bootleg and remix versions of pop tunes is mainly due to the mass availability of cheap, easy to use audio-editing software.

Richard Bown goes one better with this collection of free Linux audio apps.

There are many audio editors available for Linux and here we examine just a few of them.

An audio editor is a piece of software designed to edit sound files. These files could be generated from many sources, a microphone or instrument, a CD or even downloaded from the Internet. Audio files come in many formats but probably the most popular and common is Microsoft's WAV format. Stereo 16-bit 44.1KHz WAV is CD quality, and WAV files of this type can be written to CD without

any resampling. WAV is the *lingua-franca* of audio editor formats and the format that all editors should attempt to support as a minimum.


The basic job of an audio editor is to display a waveform generated from the audio file and then let the user edit it either by cutting and pasting sections, trimming, fading in and fading out portions, converting bitrates (known as down-sampling and up-sampling), changing from stereo to mono or vice versa (changing the number of channels) and applying

more advanced effects usually through a plugin mechanism.


Using an audio editor should be equated to using a text editor or word processor. The editor doesn't need to be fancy or complicated to do a good job, and sometimes the fewer 'killer features' on the editor the better. Some useful features to look out for are normalisation, fade outs and fade ins (useful for "top and tailing" audio segments) and occasionally useful features like hiss reduction or crackle removal (useful when recording from

vinyl). Normalisation is the process of adjusting the volume levels of an audio file to a preset maximum – this is invaluable (along with Compression and Limiting) when trying to match levels on different audio tracks (for example when preparing audio files for a compilation CD).

The reference audio editors under Windows are applications such as *Cool Edit Pro*, *Wavelab* and *Soundforge*. These are great applications which always draw comparisons. They also exemplify applications that work well as editors without being too complicated – *Cool Edit Pro* in particular succeeds in this regard and is a great benchmark against which to judge other editors.



Highlight an area of the audio file and then apply a plugin effect to it.



Split a stereo audio file by channel – you can view the amplitude of the audio file linearly or in dB.

Audacity

The workhorse of the Linux audio world – and it works on Windows and Mac machines too!

■ **VERSION** 1.1.0 ■ **WEB** <http://audacity.sourceforge.net/>

Audacity is truly cross-platform

Audacity is a very successful open audio editor based on the wxWindows GUI toolkits, which makes the GUI portable. It's available for Linux/UNIX, Windows and Mac OS 9/X.

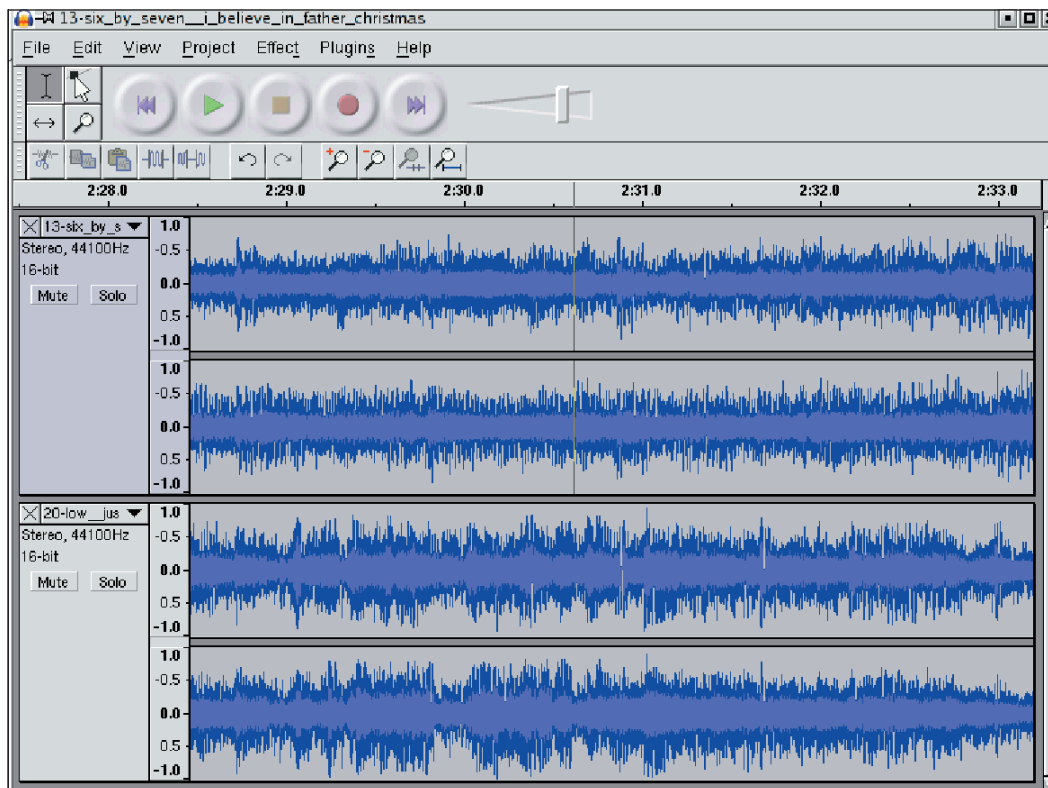
Audacity is a very successful open source project and it's a mature and well-rounded piece of software with good documentation and a large user base. It's a popular download and

offers a good set of basic editing features and a large set of LADSPA and built-in plugins. Due to its capabilities, it's pretty much become the audio workhorse of the Linux world, and has got there because it's quick, easy, looks good and has all the right features. There's multi-level undo, it handles large files with ease (even many minutes of CD quality audio), allows editing of MP3 files, and of course supports the native formats of all of its platforms. It also supports native plugin formats including VST plugins significantly.

The version under test wasn't the latest bleeding edge version, but the recommended version 1.1. Audacity uses its own project file arrangement and allows the user to import audio, raw data or even MIDI into the project on new tracks. This is similar to the *Cool Edit Pro* approach for those users wanting a little more than basic editing from their audio editor.

Each track in the project has a set of attributes to quickly tell the user the status of the audio file – whether the track is mono or stereo, what the sampling rate is and the bit depth. Tracks can be muted or soloed and the entire project mixed down (exported) to a single audio file as the finished product. As mentioned the editing features are good – the audio waveform is zoomable and selection and editing of parts of the audio is very simple.

Audacity is the undoubtedly the reference audio editor for Linux and it's the standard by which all other editors must be judged. It offers the complete package (itself somewhat of an achievement) and it's stable and reliable. Always worth checking out.



Load several audio files into a single Audacity project and mix then down to a single stereo audio file for export.

VERDICT

Features	8/10
Ease of use	8/10
Documentation	8/10
Performance	8/10

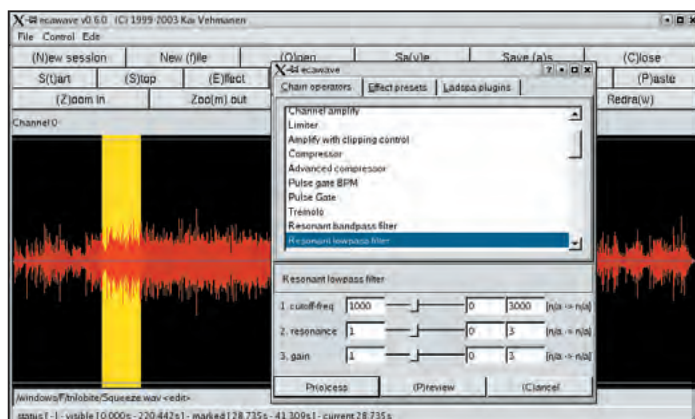
Easily the pick of the audio editors on trial here. A good solid performer with all the necessary features and some good documentation.

LINUX FORMAT RATING
8/10

Ecawave (with ecasound 2.2.0)

How easy is it to edit anything without an undo function?

■ **VERSION** 0.6.0 ■ **WEB** <http://www.wakkanet.fi/~kaiv/ecawave/>



Applying plugin effects in *Ecawave*. Be prepared to wait a while though...

Ecawave is part of the ECA project, which aims to provide a set of useful tools for audio editing under Linux. While the drive of this project is admirable, to say that *Ecawave* itself

has a basic user interface would be to stretch the point a little. *Ecawave* works with the *ecasound* library so you have to download and build both of these parts of the project to get up and running.

When you start up *Ecawave*, you're confronted with a set of buttons in a strange paved motif across the window. Opening an audio file is pretty straightforward, but immediately you notice the lack of niceties; for example, the open file dialog doesn't remember where you picked up the last file, and there's no numbering of the time axis on the waveform.

Zooming works, but is very basic. Selection of sections works pretty quickly, but applying effects is a bit time-consuming and loading files is generally a pretty slow process too. *Ecawave* doesn't have any undo facility, which will mean you'll constantly have to save your progress incrementally, but does have support for LADSPA plugins. The choice of hotkeys for starting and stopping playback of the audio file (**Control+S** and **Control+T**) are strange and not particularly obvious or easy to reach with one hand.

While *Ecawave*'s basic features all seem to be there, the user interface needs a lot of work to attract new users. The niceties are what make people come back to any editor, and this can be hostile editing environment. A lack of undo is a major flaw and some of the operations take a bit too long to accomplish.

VERDICT

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

Based on a solid architecture but sorely lacks a few crucial features that are standard in other editors, and would benefit from some UI niceties.


LINUX FORMAT RATING

4/10

GNoise

A wave file editor for GTK+ or GNOME – the ongoing personal project of Dwight Engen

■ **VERSION** v0.1.15 ■ **WEB** <http://gnoise.sourceforge.net/>



Editing of a WAV file with *GNoise* – make sure you work from copies!

GNoise seems to be pretty much a one-man project and has been in development for the last couple of years. Building it from source is a little bit of an odd experience as the bundled configure script was a bit broken with this version.

GNoise is a GNOME-based wave file editor for Linux – Microsoft WAV files only that is. It's a pretty small application, lightweight and with its major concerns being "speedy and be able to handle big files". It's reported to work with files bigger than 500MB and seems to be pretty stable even at large file sizes.

It lists its features as being able to load and display files (generating a display cache), playback, recording cut, copy, paste, (unlimited) undo, mute, fade in/out, reverse, normalize, and LADSPA plugins.

One quite major limitation with *GNoise* is that editing operations are performed directly on the original soundfile itself. Even if you don't explicitly save the file, then the edits will be performed on it anyway. For the majority of casual users who are just playing around with *GNoise*, it could come as quite a shock to find out that they've already modified all of their files without even realising it.

While *GNoise* has great intentions and has been proven to work quickly with even large files, it should have been designed to be more careful with its users files. The current type of destructive editing is pretty dangerous. The user interface is a little basic and similar to (but not as ugly as) *Ecawave*. The undo is good and works, but the operations themselves shouldn't have been constructed to alter the original file without the user's permission.

A new internal architecture is promised for the upcoming "rewrite from scratch" GTK2 version, which will use a generic back end to get buffers of sample frames; and enable *libsndfile* or *gststream* for plugin.

VERDICT

Features	4/10
Ease of use	6/10
Documentation	4/10
Performance	3/10

Appears to be a quick and powerful WAV-only editor, but in-place editing makes it too dangerous for novices.

LINUX FORMAT RATING

4/10

GLAME

Glame it on the boogie with the unique filter network system from some not-so-boring Germans...

■ **VERSION** 0.6.4 ■ **WEB** <http://glame.sourceforge.net/>

This project derives its name from the rather unwieldy, crow-barred phrase "GNU/Linux Audio Mechanics" and comes in two flavours – the

graphical front end (*glame*) and the command line interface (*cglame*). The front end looks smart and GNOME-like and is pretty easy to use. Audio

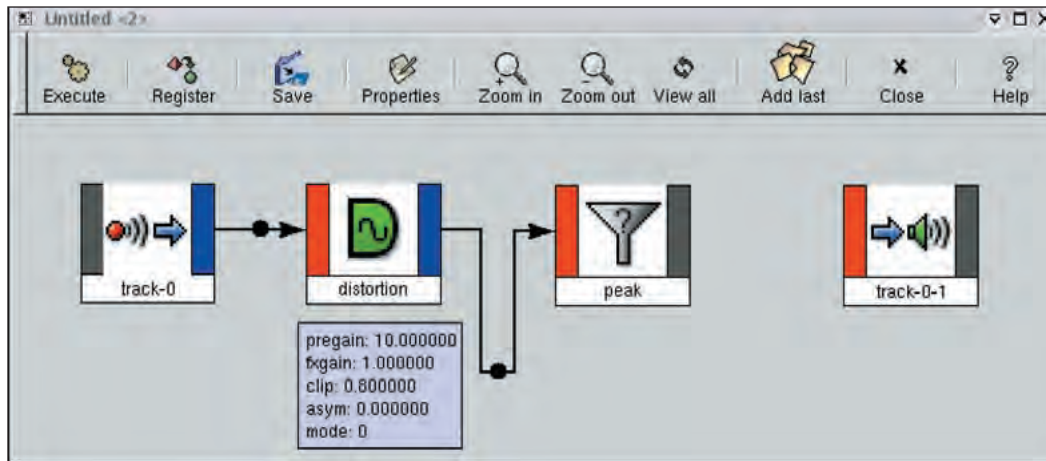
files can be imported into a project file and edited a channel at a time in smaller audio editor windows. These editor windows are quite small and appear at first glance to have little functionality but they hide a lot under frankly quite unobvious buttons and on right mouse button clicks. Editing of the audio files is pretty simple and once again the usual features are here in abundance.

Large files are handled elegantly (*GLAME* has its own swapfile concept) and files imported to projects aren't

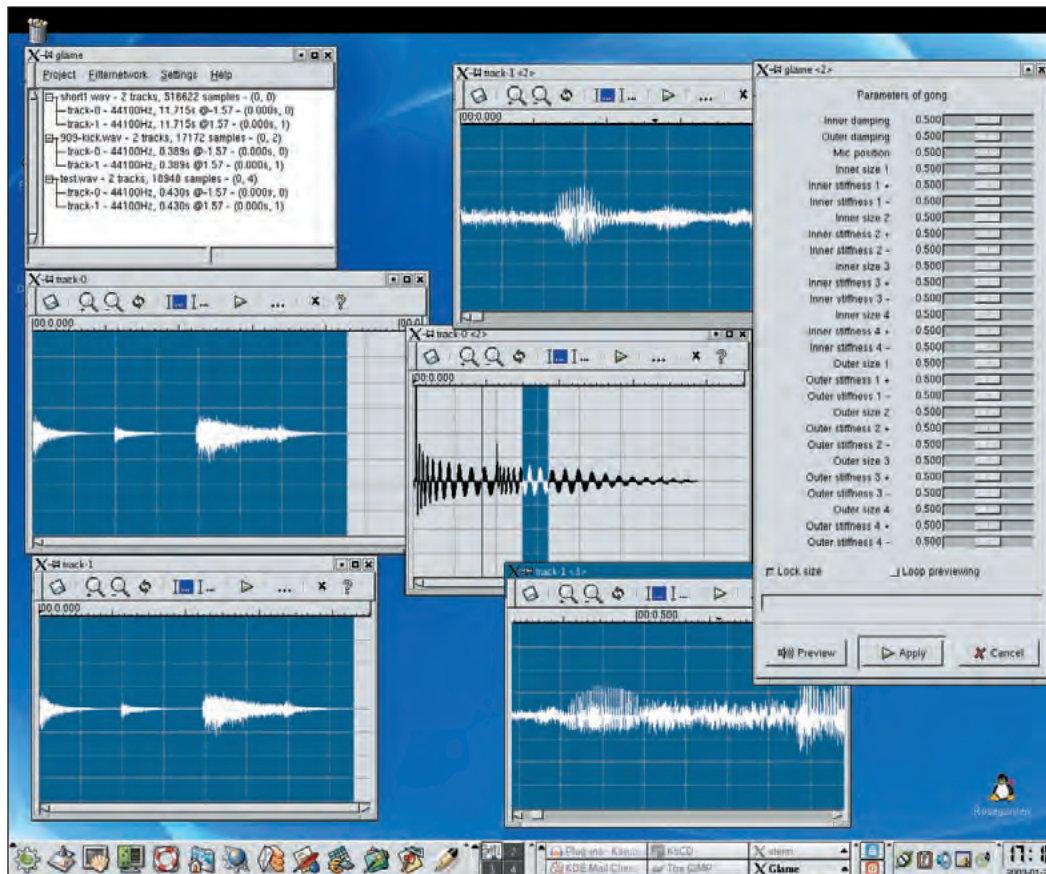
modified. The results instead can just be exported. A nice touch to *GLAME* is the filter network concept – this allows the user to create a set of filters to apply in one pass to audio files. Thus if you have a repetitive audio processing task to perform, you might find this feature pretty useful. Creation of the filter network is simple, but make sure you read the manual first. The right mouse button is used to select entries. To connect two filters, all you have to do is draw a line from the output of the first to the input of the second, dragging the mouse with the left button pressed. The name "filter network" is a little misleading too – I really didn't know what to expect when I first opened it.

GLAME suffers a bit from, chatty language in the tutorial pages which can be a bit distracting but nonetheless is quite amusing; the documentation is very well-written considering that none of the authors' first language is English. Unfortunately though, the application suffers a little from not being labelled well enough in places and not selling all of its features to the user. The user interface is neat and tidy and works pretty well (and there's undo here) but it's not always obvious how to do stuff or where to go next.

The GUI sometimes appears broken, even if it's not, because it's often really not obvious what needs to be done next. The command line interface that sits underneath the GUI has obviously been well thought out, as has the entire project, but it could do with some attention to usability testing and some more work to reduce its quirkiness and make it more appealing to novices. While it's nice to have apps that work in an idiosyncratic style, it's usually better for them to be boring and just be good at their job. *GLAME* is good at its job but its sometimes difficult to work out how to use it.



GLAME filter networks. Create your own custom filter sets and apply them to your audio files.



GLAME and its many windows. Here we're applying an effect to a section of audio file.

VERDICT

Features	7/10
Ease of use	5/10
Documentation	6/10
Performance	7/10

Attractive and well thought-out, with real depth and imagination, but the idiosyncratic UI might not be to everyone's taste. Filter networks feature is innovative and exciting.

LINUX FORMAT RATING
7/10

RoundupAudioEditors

Sweep

Not Sooty's cheeky doggie chum, but a capable performance application that edits as well.

■ **VERSION** 0.8.0

■ **WEB** <http://www.metadecks.org/software/sweep/>

To pardon the obvious pun, Sweep

is making a lot of noise in the Linux audio world. *Sweep* has been sponsored by PIXAR (which itself is pretty exciting) and the main author, Conrad Parker, is quite excited about what he's created and rightly so.

Sweep is mainly aimed at the demands of live performance of audio but still provides good editing features and it certainly looks like a polished and professional application.

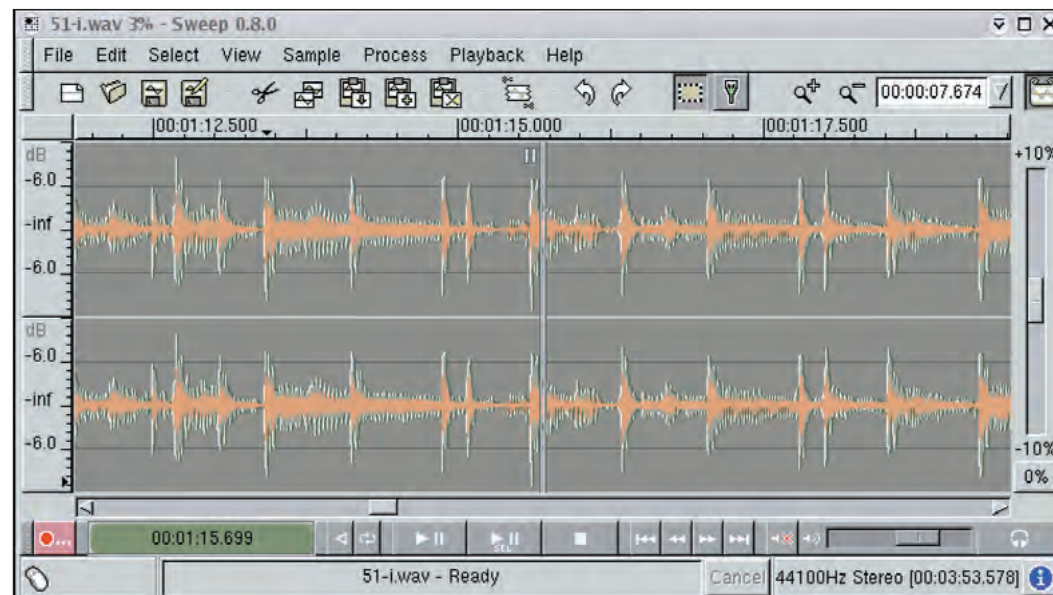
From the documentation "[*Sweep*] supports many music and voice formats

including WAV, AIFF, Ogg Vorbis, Speex and MP3, with multichannel editing and LADSPA effects plugins." Much is made of the stylus scrubbing tool called "Scrubby" which allows the user to scratch audio files in real-time just like you would a record on a turntable, and it's indeed a pretty nice tool and a great idea, but it probably takes someone with a modicum of practical skill and experience to get the most out of this feature. Real-time audio effects are plentiful – reverse playback, playback speed can be modified in real-time just like you would with a DJ turntable. Recording into segments of an audio file is easy and is again attractively and clearly presented.

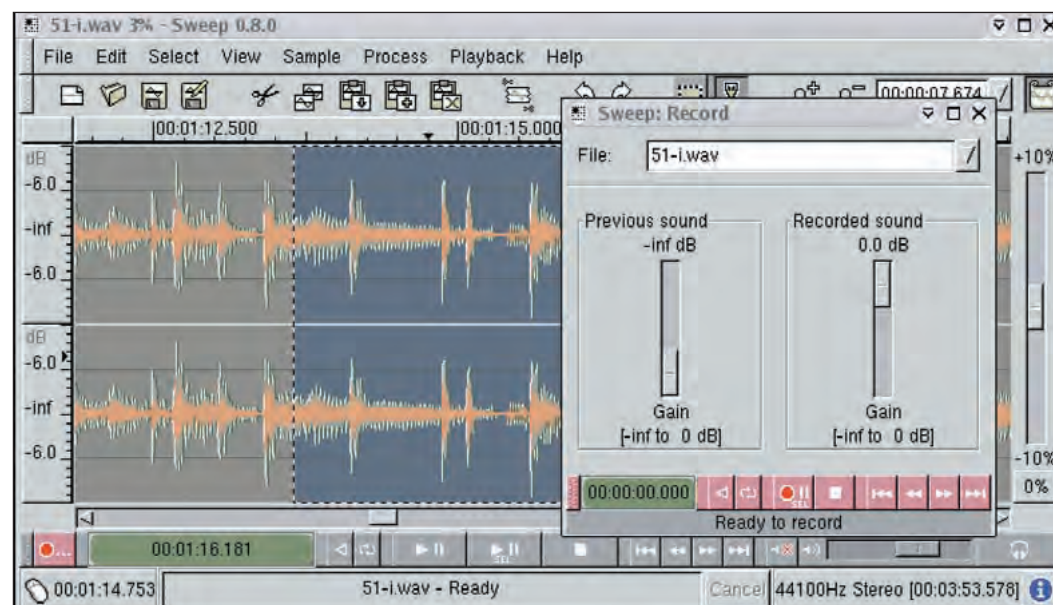
Additional highlights are facilities for looped and reverse recording, internationalisation, Ogg Vorbis, MP3 and Speex compressed formats, multiple views, discontinuous selections and unlimited undo/redo.

Sweep is a very polished product and Conrad Parker makes sure that any new features are necessary and complete before they go in – an example of bloat-busting that quite a few Open Source developers could do well to emulate. The audio editor looks good and feels good to use too but it's tipped slightly in favour of audio performance than it is to editing. There's no use of disk caching when editing, so large files aren't handled very well (*Sweep* will attempt to load them all into memory, which can cause a few problems). Also while operations like fade in and fade out and normalize are provided they're located amongst other operations in a complex menu hierarchy. Common operations like these should be conveniently located for easy access. This is a minor inconvenience though, and will undoubtedly be addressed in future iterations. On a personal note, I'd also like to see a 'file preview' option on the audio file browser too, and I actually went as far as submitting a patch to do this – it's not made it into the distributed code yet, but a feature like this on an audio editor is very useful feature if you're browsing a large selection of audio files trying to select a tune to work with.

Sweep is a great looking, full featured audio editor and great fun to play with. It certainly looks better than its main rival, *Audacity*, but it does tend to favour performance over functionality as an audio editor. It will continue to lead the pack alongside *Audacity* but some effort might have to be spent improving the audio editing capabilities.



***Sweep* is a great looking and full of features, especially for performing real-time effects on audio files. If you want to play to a crowd without the hassle of humping large boxes of vinyl around, *Sweep* is a DJ's best friend.**



***Sweep*'s record dialog make it simple to record new audio into specific sections of audio files.**

VERDICT

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

Not only a good audio editing tool (and support for many audio formats) but also great performance options, tweakable playback speed and great scratching tool.

LINUX FORMAT RATING
8/10

mhWaveEdit

Basic taken to the point of spartan...

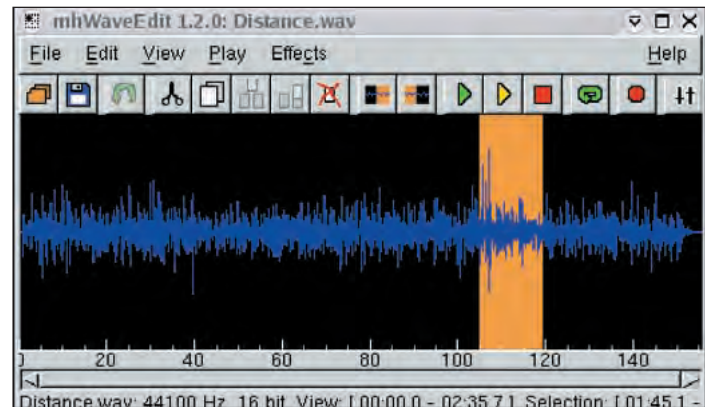
■ **VERSION** 1.2.0

■ **WEB** <http://www.mtek.chalmers.se/~hjormagn/mhwaveedit.html>

mhWaveEdit is another one-man project (written by Magnus Hjorth – presumably hence *mhWaveEdit*). It loads, plays, records and saves WAV files “and a few other formats”. Editing is performed in memory if the file is small, otherwise it’s on disk. There’s zooming on the main window, selection is easy and straightforward, but effects are a little thin on the ground. There’s no support for

LADSPA effects and the built in effects are basic – resampling, mono to stereo and vice versa, inserting silence but no fade in or fade out and there’s also no normalisation.

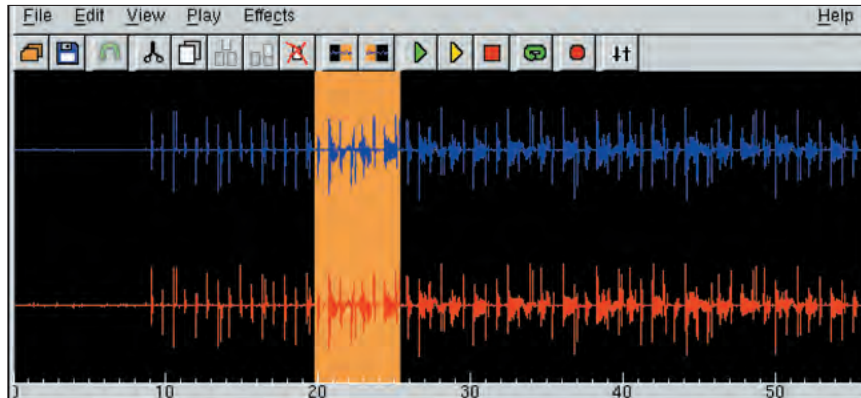
At least *mhWaveEdit* “never changes your files behind your back” but the editing functions just really aren’t there. It’s fine if you don’t want to do anything complicated with your file at all, but even the simple task of



mhWaveEdit's main editing screen. Highlight an area of audio file and apply an effect. You'll find this a lot easier using the Zoom function.

topping and tailing audio files ready for CD to play on your living room stereo isn't possible. While the project looks

neat and has potential, it should really aim to boast a few more features to attract serious users in later versions.



Combine channels – one of mhWaveEdit's few functions.

VERDICT

Features	2/10
Ease of use	6/10
Documentation	2/10
Performance	4/10

A promising editor as it stands so far, but lacks features. Future iterations could be worth looking out for if they provide greater levels of functionality.

LINUX FORMAT RATING

4/10

AUDIO EDITORS THE VERDICT

Covering six audio editors we have seen quite contrasting approaches – from the traditional user interface of *Audacity* and *mhWaveEdit*, to the less conventional look and feel of *GLAME*. We've also seen a range of functionality on display and how omitting critical features or design shortcuts can affect usability quite dramatically. It's encouraging to see that after many years in the relative

wilderness of conflicting and half-completed sound drivers that there are now starting to emerge some truly excellent audio projects for Linux that “just work” from a compatibility point of view. Some of the features on offer in these pieces of software rival (and in some cases better) their proprietary Windows and Mac counterparts and of course the fact that they are all free makes it even better news for anyone wanting to try them out.

Two projects stand out above everything in this review – these are clearly *Audacity* and *Sweep*. *Audacity* is the perfect cross-platform audio editor. It's well put together and well supported and goes about its job in a no-nonsense fashion. *Sweep* is a bit flashier and showier and is more slanted at audio performance than editing, but still does a great job and certainly looks the part. For editing audio, you can use either but probably

favour *Audacity* in its simple layout and approach – it's the closest to *Cool Edit Pro* and that's never a bad thing to be. For live audio performances then the choice must be *Sweep*. It's good that there are two audio products of such a high now standard available for Linux. It surely can't be long before projects such as these (in concert with other Linux audio projects) can take Linux audio into the realm of professional music-making and beyond. [LXF](#)

Table of features

Name	Version	Sound Drivers	Formats	Requirements	Comment
Audacity	1.1.0	OSS	Most common	wxWindows	Full-featured and stable, well documented.
Ecawave	0.6.0	OSS/ALSA	Most common	ecasound	Basic and plain with a lack of features.
GLAME	0.6.4	ALSA	WAV	GNOME libs	Interesting design and many useful features – not least filter networks.
GNoise	0.1.15	OSS	WAV	GNOME libs	Able to handle large (WAV only) files but destructive editing could be dangerous.
mhWaveEdit	1.2.0	OSS	WAV	SDL libraries	Neat and simple but too few features.
Sweep	0.8.0	OSS/ALSA	Most common	GTK libraries	Beautiful look and feel, tends to favour performance over editing features.

HotPicks

The best new open source software on the planet!



Richard Drummond searches the Net and badgers developers for new apps to review for LXF.

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

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

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

HotPicks at a glance

DNSMasq	40
Trackballs	41
distcc	42
KVocabulary	42
Seahorse	43
Pixie Plus	44

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!



DNS PROXY

DNSMasq

■ **VERSION** 1.11 ■ **WEB** <http://thekelleys.org.uk/dnsmasq/doc.html>

```
Shell - Konsole
File Edit Search Buffers Windows System Help
# Configuration file for dnsmasq.
#
# Note that this file supercedes /etc/default/dnsmasq which
# still exists but whose use is deprecated for most settings.
#
# Format is one option per line, legal options are the same
# as the long options legal on the command line. See
# "/usr/sbin/dnsmasq --help" or "man 8 dnsmasq" for details.
#
# Change these lines if you want dnsmasq to serve MX records.
# Only one of mx-host and mx-target need be set, the other defaults
# to the name of the host running dnsmasq.
#mx-host=
#mx-target=
#selfmx
#
# Change this line if you want dns to get its upstream servers from
# somewhere other than /etc/resolv.conf
#resolv-file=/etc/ppp/resolv.conf
#
# Add other name servers here, with domain specs if they are for
# non-public domains.
#server=/localnet/192.168.0.1
#
# You no longer (as of version 1.7) need to set this to enable
# dnsmasq to read /etc/ppp/resolv.conf since dnsmasq now uses the
# "dip" group to achieve this.
#user=
#
# If you want dnsmasq to listen for requests on only one interface
# (and the loopback) give the name of the interface (eg eth0) here
#interface=
#
# Change these if you want dnsmasq to cache any "hostname"
# or "client-hostname" from a dhcp's lease file
#dhcp-lease=/var/lib/dhcp3/dhcpd.leases
```

Sharing your dial-up server's DNS settings with a local network is made easy with DNSMasq.

Sharing an Internet connection between machines on a local network is easy with a Linux box set up for masquerading – it's just a matter of concocting your *iptables* rules correctly or using an off-the-shelf solution to do the job. Most Internet protocols should then work transparently on your client machines. What's less straightforward to sort out – particularly if you have a dial-up or cable Internet connection – is dealing with DNS. Your masquerading server will get its DNS settings via PPP or DHCP, but what about the other machines on your network? You can statically configure them to use your ISP's name servers, but this is a less than elegant solution. What happens when you are off-line or when your ISP's DNS set-up changes? Well, help

is at hand with *DNSMasq*. It's a lightweight DNS proxy designed just for these types of situation.

DNSMasq runs on your masquerading server and forwards DNS request to your ISP's DNS servers upstream and caches the results. You simply tell your local machines to use your masquerading server for DNS, and Bob's your father's brother. Actually, there's a lot more to it than this, but in essence using *DNSMasq* is really that simple. What makes *DNSMasq* truly useful is the other tricks it has up its sleeve.

For one, you can either specify the addresses of your network's upstream name servers statically or *DNSMasq* can read them from one or more named files. By default it gets the DNS servers from */etc/resolv.conf*, but you could

specify */etc/ppp/resolv.conf* if you're on a PPP connection using dynamic DNS or */etc/dhcpd/resolv.conf* if you get your network settings via DHCP. (You can even specify both files, and *DNSMasq* will pick and use the most up-to-date one. This is great for laptops where you might want to use a dial-up connection when you're on the road, or a local network connection when you're in the office. Proof that *DNSMasq*'s usefulness isn't restricted to servers.)

Another feature of *DNSMasq* is that it can serve the names of local machines via DNS too. For instance, by default it will serve up the contents of your */etc/hosts* file. But it can also be used in conjunction with a DHCP server. It can look at the */var/dhcp/dhcpd.leases* file that your DHCP server creates and add any named leases it finds there to its table of known hosts. This means that, if your local machines provide a host name when they ask for a DHCP lease, *DNSMasq* can associate the IP address of their leases with their hostnames and so provide them via DNS to the rest of the machines on your network.

Ease of use

Are you convinced yet? If not, *DNSMasq* has more features than we have space to talk about here. For instance, it can serve up MX records, which, for example, could point to a mail relay on your local network; it can forward DNS request to different servers based on the domain name requested – useful for accessing additional private name servers; and it understands IPv6.

I've been using *DNSMasq* on my own network for a couple of years now, and can't imagine life without it. It's simple to install and use, and has complete and digestible documentation. The latest release can read its settings from a configuration file rather than simply its environment, and so is easier to use than ever. *DNSMasq* can dramatically reduce the amount of administration required for networks which use dynamic DNS, and hence is an indispensable tool for those with such networks.

GAMING CLASSIC

Trackballs


■ VERSION 0.4.1 ■ WEB <http://www.lysator.liu.se/~mbrx/trackballs/>

There are people that criticize the Open Source movement for a lack of innovation. Often this complaint is justified, particularly when it comes to games. After all, how many lacklustre clones of *Tetris* does one really need? But this lack of originality can be a good thing too – it's great to some of the gaming gems of the past re-implemented and updated for modern platforms and hence gaining the benefit of 3D graphics and 16-bit audio. *Trackballs* is just such a case. It's an unabashed copy of that grand old Amiga classic, *Marble Madness*.

The idea of *Trackballs* is devastatingly simple. Guide your marble around a maze of platforms and ramps, avoiding hostile marbles, spikes and acid, and negotiating precipices and get to the exit within the allotted time. Points are awarded for picking up flags along the way, bumping off enemy marbles and for beating the time limit. Controlling your marble is a constant battle against gravity and friction, and life is made more difficult by the steepness of slopes and a variety of different surface types such as sand or ice. Picking up powerups such as spikes can give you more traction, though.

Trackballs updates on its progenitor by offering OpenGL-based 3D graphics – rather than the 'fake' isometric 3D of the original. This is used to good effect with neat touches such as light sources: game levels can be set during the day or at night. A hardware-accelerated 3D card is definitely required for players to get the best out of the game. Controlling your marble is done with the mouse, and you need a good framerate for the controls to be responsive.

As is common for GL-based games, *Trackballs* is built on the SDL library. Compiling it is straightforward, with one caveat: it requires an up-to-date version of *Guile*, the GNU Scheme interpreter. Version 1.6.0 or better is need, whereas most distros are still shipping 1.4. Still, getting hold of and building *Guile* is itself no big chore. *Guile* is commonly used for scripting plug-ins for applications, so



Right from the title screen at the start, *Trackballs* looks pretty special.

what's it used for in *Trackballs*? Glad you asked. *Trackballs*' levels, interestingly, are coded as *Guile* scripts. Don't worry if you don't parlez vous Scheme, because *Trackballs* also includes a level editor.

Trackballs oozes quality. It's still in the early stages of development, but it's quite playable and already the level of polish impresses – from the Star Wars-like credit sequence on the title page, to the professional sound track, to the in-game graphical effects. And, while it's unfinished, it's still immense fun to play, even in its present state.

Wrinkles


There's always plenty of scope for improvement. The most crucial need is for more levels. The six that are included will keep you occupied for a time – especially with the three difficulty levels – but more are definitely needed. A greater variety of puzzles and obstacles would help to expand the game, too, as would more and more varied enemies and more objects to pick up. A greater graphical diversity would be good, too. Perhaps levels could have themes – in addition to being night or day.

Half the battle to creating new levels would be over if the author were to iron out some of the wrinkles with the level editor. At the moment, it's interface is a little idiosyncratic to say the least – some documentation would probably be useful. Document the key strokes, add mouse control, maybe even a rudimentary GUI and Joe Public (that's all of us) could help out with the level design.


Finally, finishing touches such as being able to edit game settings within



Avoid the obstacles, such as spikes and acid, and collect the flags.



The level editor shows potential but the interface can be confusing.



It all looks downhill from here – a chance to show off the game's physics.

the game should be added. Currently, controls such as choosing a screen mode, picking a level to start on and launching the level editor must be done with switches via the command line. A 'Save Game' facility would be useful, too, or at least the ability to carry on at the level you last played.

Trackballs has lots of potential to grow into a really first-rate game, worthy even of the classic status afforded to the original *Marble Madness*. Even if it takes some time to get there, right now it's a well-presented, addictive piece of gaming which is well worth trying out.

COMPILER TOOL

distcc

■ **VERSION** 1.1 ■ **WEB** <http://distcc.samba.org/>

The ability that open source software gives you to compile and build software yourself offers tremendous advantages, but often the amount of time it takes to compile any non-trivial piece of software is daunting. If you have a local network of machines handy, then you can significantly reduce this time with *distcc*.

distcc facilitates the parallel building of software by offloading compilation jobs to machines on the

network. It's deceptively simple to use. Each machine on your network runs a *distcc* server and the compile host runs the *distcc* client. This client acts as a front-end for *gcc* (or *g++*). It launches *gcc* to do any preprocessing required, and then sends the preprocessed file over the network to a *distcc* server for compilation, which, when finished, returns the compiled object file. Since only the task of turning preprocessed source code into object code is done remotely, it is only

necessary that the *distcc* servers have a compatible version of *gcc* installed – they don't require any of the header files or link libraries used.

The task of parallelizing the build is actually done by *make*. You pass the **-j** option to *make* to specify the number of *make* jobs to run concurrently. When you are using *distcc*, extra jobs are off-loaded using a simple load-balancing heuristic.

distcc is compatible with the auto tools, *gcc* and *g++*. It even copes with cross-compiling, as long as you have a correctly installed cross-compiler on each of your *distcc* servers. *distcc* launches the compiler binary with the same name on each *distcc* server, so, for example, if you were building for a Linux PowerPC target, you would specify your compiler (to your *make* or

configure scripts) as

CC=distcc powerpc-linux-gcc

If one of the hosts is a PPC box, you would probably have to ensure that a link in fact exists between *powerpc-linux-gcc* and the native compiler, *gcc*.

The performance win offered by *distcc* depends on the task at hand. While it won't scale linearly, high degrees of efficiency are possible for projects with enough potential parallelism to exploit. The Linux kernel, for example, benefits greatly from parallel building with *distcc*. Even if you don't have a big network of machines to use, *distcc* is useful for borrowing the resources of a more powerful machine when building on a slow machine such as laptop.

VOCABULARY TRAINER

KVocabulary

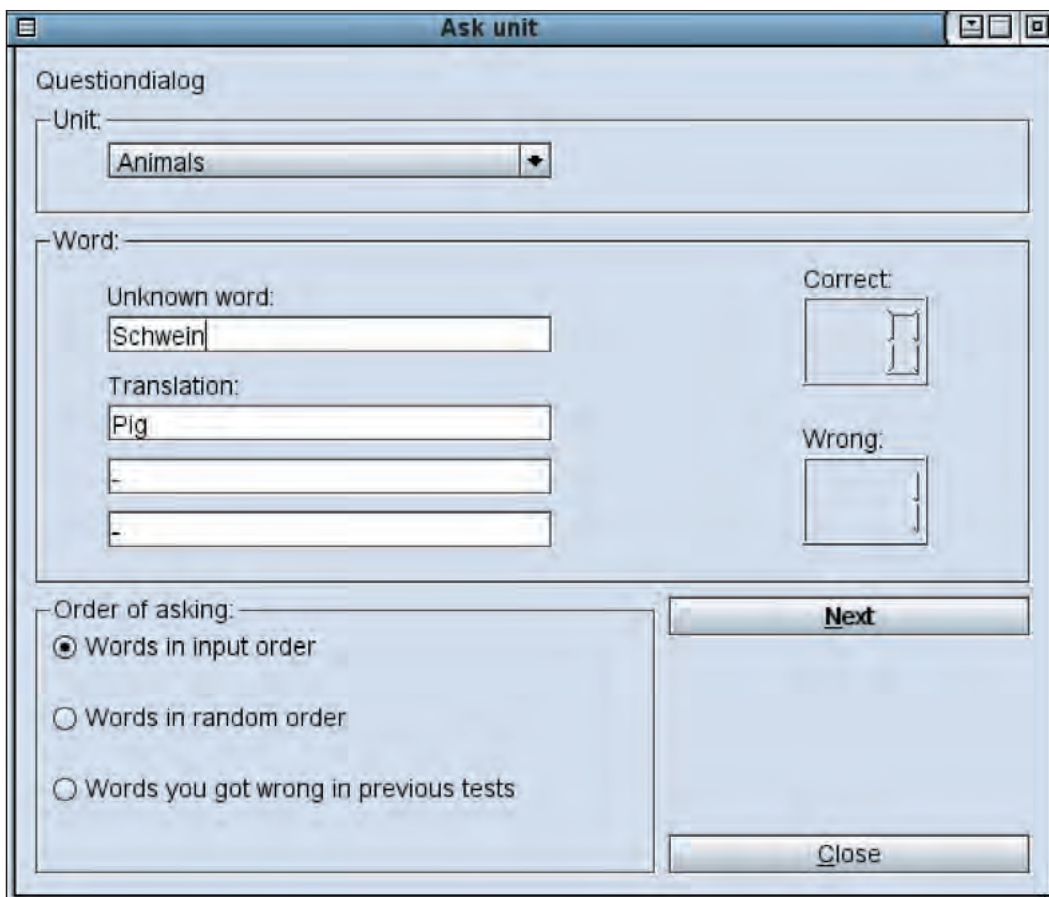
■ **VERSION** 0.2.7 ■ **WEB** <http://linux.handorf-langenberg.de/>

Learning the vocabulary of a foreign language can be a very time-consuming endeavour. Often the best approach is simply repetition. *KVocabulary* is a little KDE application that can help out here. If you're having to spend a lot of time in front of your

monitor anyway, why not while away those odd moments by brushing up on your vocab for your foreign holidays rather than just staring into space?

KVocabulary is very rudimentary. It lets you compile your own lexicons of words to test yourself. It supports any number of lexicons and words in any of the languages supported by KDE. Each lexicon is subdivided into units or topics. To build up a lexicon, you enter the foreign word – the word to be tested – and its translation (the translation can have up to three lines). In test mode, or what *KVocabulary* calls 'ask' mode, for each word in your selected lexicon it will show you the translation and ask you to enter the foreign or unknown word. It can either ask you them in order or randomly, or simply ask you those that you got wrong in previous attempts.

If you need a hand learning vocabulary, then perhaps *KVocabulary* will help, but it could do with some improvements. I've always found in languages that I've tried to learn that it's the gender of nouns and the various inflections on words depending on case that are difficult to remember, but, currently, *KVocabulary* offers no real support for helping this kind of learning. The GUI needs an overhaul as well. At the moment its layout is none too clear. For example, when it 'asks' you a word, it would be clearer if it actually posed you a question, rather than just presenting you with some string widgets. But then, who said learning another language was easy?

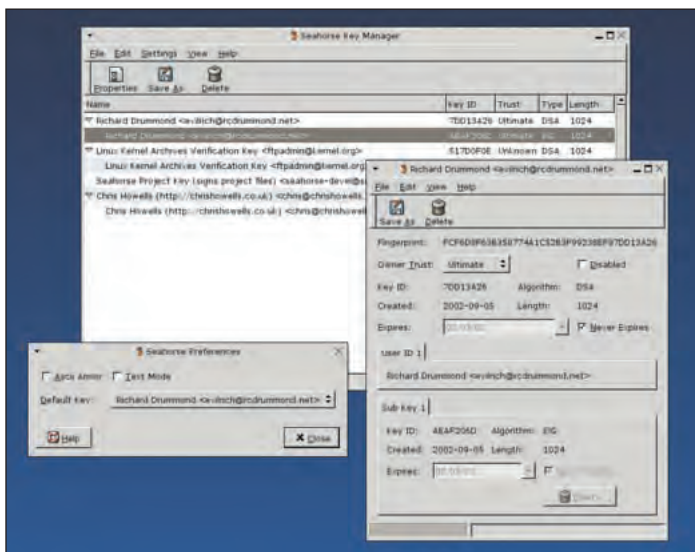


KVocabulary is a simple tool to help you learn the vocabulary of another language.

GNUPG FRONT-END

Seahorse

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



Seahorse's key manger simplifies the manipulation of GnuPG keys.

You don't have to paranoid about snooping third parties eavesdropping on your email conversations to use a personal security tool such as *GnuPG* (www.gnupg.org). *GnuPG* is the open source clone of Zimmermann's famous *PGP* – Pretty Good Privacy – and it has more uses than just simply encrypting files. In fact, the most significant application of *PGP* and hence *GnuPG* is the creation and verification of digital signatures. They let you tell whether a file was actually created by its claimed author and whether it has been tampered with en route to your machine. Digital signatures – for examples on software distributed in RPM format – form an important defence against viruses.

The problem with *GnuPG* is it's a less-than friendly tool to use. It's strictly command line-based, and using it to perform any of the tasks of which it is capable involves grappling with the requisite series of switches and options. Personal security isn't much good, after all if you can't use it properly. Luckily, there are several good graphical front-ends to *GnuPG*, which can ease the whole process. One such is *Seahorse*, a GNOME-based tool, which has just been ported to GNOME 2.0.

Seahorse isn't the most powerful of *GnuPG* front-ends; its rival *GPA* offers

more functionality. But *Seahorse* is a lot easier to use, and boasts very comprehensive online documentation to help out the first time user.

Building *Seahorse* should present few problems if you have the various bits and pieces that make up GNOME 2.0 installed. *Seahorse* interfaces with *GnuPG* via the GPGME library (also used by *GPA*) For *GnuPG* Made Easy – see www.gnupg.org/related_software/gpgme/

The *Seahorse* interface is split into three main functions: key manager, file manager and text editor. The key manager is obviously where you create, sign, import and export keys and so on. The file manager lets you use those keys to sign files, verify the signatures on files, and to encrypt or decrypt files. The unique feature of *Seahorse*, the text editor, allows you to perform similar operations on a text buffer – without having to load or save that data to a file. This is great for cutting and pasting and hence manipulating signatures or encrypted data if, for example, your preferred mail client can't directly use *GnuPG*.

The key manager lets you create key-pairs with a straightforward dialog. It supports the creation of DSA/Eigmal keys (for both signing and encryption), DSA keys alone (for signing only) and RSA keys (for signing). You can set

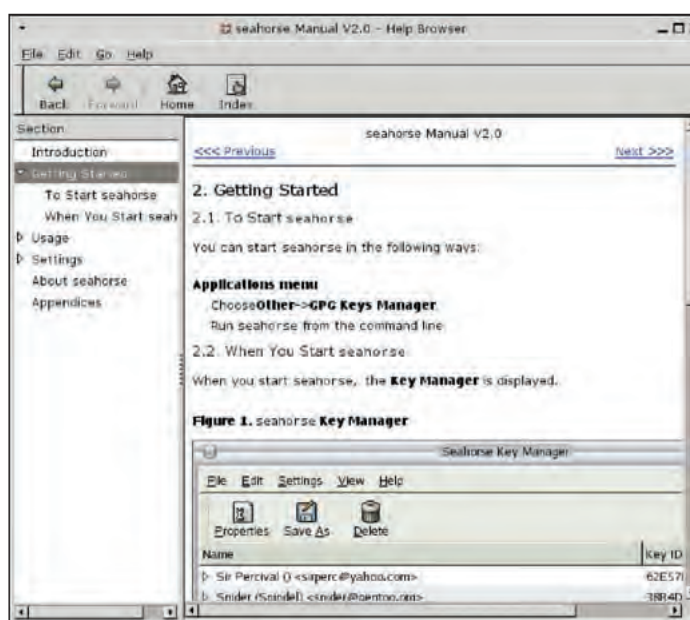
your key length, pick an expiry date, enter a pass phrase and all the usual options. You can then save out your public key as a file for distribution. *Seahorse* doesn't support the generation of a revocation certificate for keys it generates, nor can it upload keys to a key server.

In fact, *Seahorse* cannot export or import keys to or from a key server at all – yet – and this is the most serious shortcoming with the program. (Both *GPA* and the Qt-based *GPGKeys* do support key servers). The key manager also won't let you generate additional sub-keys and, while it does show the trust level of a key, it doesn't list which keys have been used to sign a key.

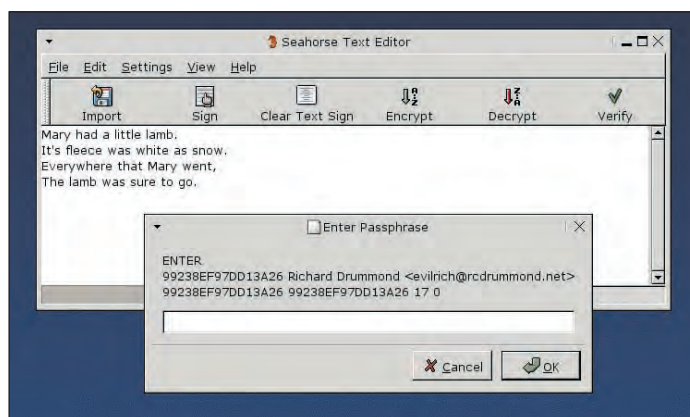
The file manager, on the other hand, offers all the features you need and is intuitive to use. Signing a file is a simple matter of picking a file in the file dialog, click 'Sign' and entering your

pass-phrase. Verification is even-easier: just pick the signature file and hit 'Verify'. Encryption is equally easy, but this time a dialog pops up where you can pick the keys of the intended recipients of the file. The file manager also lets you export your key chain to a file or import multiple keys from a file.

Seahorse is great for those new to *GnuPG*. It presents most of the functionality you need in a clear and well-presented interface. Missing features, such as no support for key servers, might be a worry, and it would be nice if *Seahorse* could remember your default key's passphrase during sessions to stop you having to input it over and over again. Its integration with your desktop environment could also do with being improved, such as, for instance, allowing you to drag-and-drop files to be signed, verified, encrypted or decrypted.



The online documentation, presented by GNOME's *Yelp*, means that *Seahorse* is easy to learn for those wanting to use *GnuPG* with a GUI.



Signing, verifying, encrypting and decrypting files is a doddle with the file manager, but we'll have to wait for proper drag-and-drop functionality.

IMAGE PROCESSOR

Pixie Plus



■ VERSION 0.5.3 ■ WEB <http://www.mosfet.org/pixie/>



Pixie lets you view, manipulate, and process batches of images quickly.

The motto of the Perl community, 'There's more than one way to do it', is in fact more than appropriate for the Open Source movement as a whole. The sheer number of open source projects gives the end user an unparalleled degree of choice. Take for example, image processing. I wouldn't like to sit down and count the number of tools available for munging and manipulating images in various ways. Even if you specifically want a KDE application for the job, there are alternatives aplenty. The official KDE tree alone offers – in order of power – *KView*, *Kuickshow*, and *Krita*. If none of these fit your bill, then why not give *Pixie Plus* a whirl?

Pixie Plus is fast and feature-packed image viewer and processor that is particularly suited to handling large quantities of images. It boasts a range of image processing effects which it can apply to batches of images, and *Pixie* has its own optimized image I/O system, based on *ImageMagick*, which supports up to 80 different image formats.

The main *Pixie* window is split into three areas: a familiar-looking directory tree, a thumbnail view of the current directory, and a preview of the selected image in the current directory. These can all be dragged around and positioned within the window as suits you best. This main interface works like

a file manager – with added features for manipulating images. You can do file operations such as renaming, modifying properties and copying or moving files. *Pixie* integrates well with the KDE desktop and you can drag-and-drop seamlessly between *Pixie* and *Konqueror*.

The thumbnail view, as its names suggests, lists the contents of the current directory with any images rendered as thumbnails. *Pixie* can either use KDE's thumbnail system or its own internal system. The latter is fast and supports embedded thumbnails in TIFF and JPEG files. In testing, I couldn't get either to work – images were presented merely as icons, just like other files – but this is probably due to my unstable CVS version of KDE 3.1.

Sort your image files

You have various options for sorting images in the thumbnail view. You can sort in ascending or descending order by name, date or size; additionally you can choose to show only images and directories, to show images first or to show categorized images first.

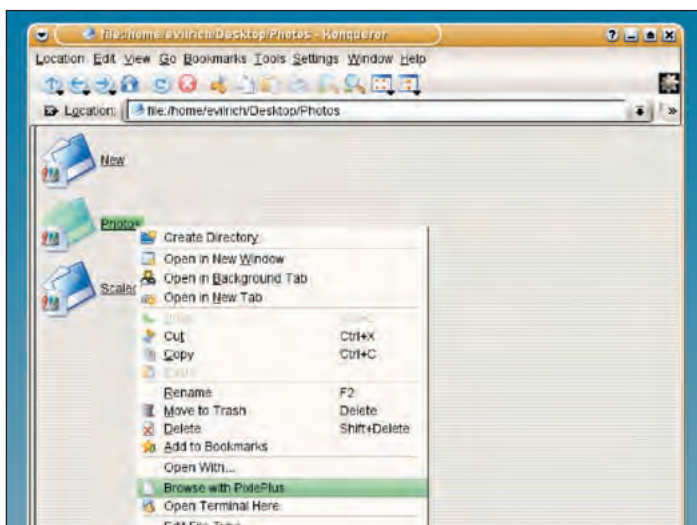
The thumbnail view provides the controls to manipulate images either singly or in batches. Here you can auto-rename and categorise images, convert the format of an image or employ any of over 30 different image processing effects – from simple rotation and scaling to blurring and sharpening to special effects such as swirl and implode. These are all based on the well-proven foundation of *ImageMagick*. When applying effects in batches, you are given the options of doing this interactively. For each

image, *Pixie* will show you the result and give the choice of accepting or rejecting that effect.

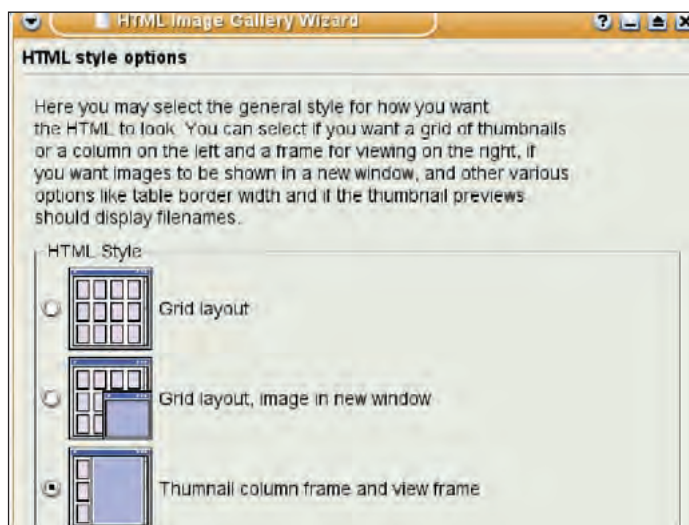
As well as an image processor, *Pixie* can be used as a straight image viewer. Double-clicking an image in the thumbnail list will open it in the viewer. *Pixie* supports various viewing modes and even has a handy full-screen slide show function. However, *Pixie* can also be used as a stand-alone image viewer. The *Pixie* distribution installs an additional *aplnk* which launches *Pixie* in this so-called express mode. The selected image will be shown in the viewer without the main image manager interface. This feature is rather spoiled by the fact that when you close the image view, the main interface pops up, though.

Another interesting and powerful features include the web gallery creator. This can generate HTML pages and thumbnails suitable for presenting your image collections on the web. A number of different page formats and customizations are supported and the results can be impressive. It turns the task of sharing your pictures on the web a real no-brainer.

For those that need to manage and process large quantities of images, perhaps scans or digital photos, then *Pixie* is hard to beat. Its interface is quirky in places, but this interface lets you work with batches of images quickly. Combined with its rendering and processing speed, this make *Pixie* a real winner. Improved support and more control over for printing would be nice, and, maybe one day, integration with *SANE* and *GPhoto* for acquiring images. [LXF](#)



Pixie Plus integrates seamlessly with the rest of the KDE desktop.



The HTML gallery generator is great for sharing your photos via the web.

Copyright

YOU, COPYRIGHT AND THE LAW



To reflect the internet age, laws of Intellectual property are changing – but in whose interests? **David Harris** investigates the copyright caper.

cover feature



Disclaimer

This article is based on UK copyright law, not US or European law. There's a big difference between the three areas. Nothing in this article is intended as or should be construed as advice or acted on without seeking your own lawyers' advice. The author cannot give personal legal advice.

As signs of the times go, it is a sad indictment of the state of the law that *Linux Format* has to run articles on intellectual property and copyright. Many of us would rather have our fun tinkering with code, hacking or arm-wrestling a distribution. My relaxation from the complexities of life as an intellectual property lawyer was to return home and hack. No more: Like others, I now have to look over my shoulder for RIAA's lawyers: Is that P2P system I'm hacking likely to result in a summons for contributory copyright infringement? Dare I download *deCSS* to watch a DVD that I legally purchased from HMV? Unfortunately, due to their need to create profit for shareholders in the short-term, corporate interests view free stuff in many of its forms as a lethal threat to their power. To take Microsoft as an example, Free Software threatens its vision of a society where all digital systems (from the phone and PDA to the PC) run Microsoft's proprietary products – a world where users are relieved of the awful burden of choice, in exchange for a relatively modest licence fee that rewards Microsoft for its 'innovation'. The content industry likewise views Free as a threat, and like Microsoft, its decision-making is dominated by a crippling paranoia over a perceived loss of control, rather than being motivated to use the Open Source model to find new profit-generating ways to improve both its products and customer service. In the digital age, it has led them to seek

ever-greater control via paid-for laws using the vehicle of copyright law. It does of course have a legal and moral right to prevent people misusing its intellectual property. However, when it uses this as an excuse to further various hidden agendas and it does so in an abusive manner, it is hardly surprising that people come to question the value of intellectual property law and its future. But before looking at where copyright law is going, let us examine from where it has come.

A brief history

Copyright law in its modern form had its genesis in the era following the invention of the printing press. Prior to that, the act of copying was so laborious and rarefied an act that it needed no regulation and caused no problems. The only people doing any significant copying were monks who would spend years transcribing, bent over vellum parchment. This work was typically an act of religious devotion or an instruction from a feudal lord, and as such the authority for it was often either God or his earthly emissary, which provided great insulation against litigation. With the invention of the printing press, the cost of reproduction dropped enormously: once printing plates were prepared, multiple identical copies could be reproduced in the time it took to spin the handle of the press.

Creating the plates still took effort and marketing incurred expense, but with literacy no longer being confined largely to priests, the larger market

made a business feasible for entrepreneurs. The entrepreneurs, called stationers, did not take kindly to unauthorised reproductions and they formed a trade guild called the Stationers Company. It introduced rigid controls over authors and reproduction of works; and its views formed the trade custom on which subsequent laws were based. It petitioned for powers and laws from Henry VIII, who was at this time looking to undermine the power of the Catholic Church and was keen to bring printing under control. He did so by banning book printing other than by the Stationers Company and by outlawing importation of books. Stationers acquired the power to seize infringing books, issue fines and the perpetual right to print all books. At this point the compulsion to draw analogies with the RIAA and the legislative heirs of Henry VIII is almost irresistible, if only to demonstrate 500 years of consistent motives and strategy by entrepreneurs. I leave that as an exercise for the reader.

After a brief period of legislative turbulence, the World's first copyright statute was passed: the 1710 *Statute of Anne*. This spoke of the effect on authors and stationers of the piracy of books as being "to the very great detriment and too often to the ruin of them and their families". It also justified itself by promoting the Act as "for the benefit of society as a whole by promoting learning" and that its purpose was "for the Encouragement of Learned Men to Compose and Write useful Books". The Act recognised that this monopoly power might have the effect of distorting price, and provided that where a stationer sold "at such a

Price or Rate as shall be Conceived by any Person or Persons to be High and Unreasonable; It shall and may be Lawful for any Person or Persons to make Complaint thereof to the Lord Archbishop of Canterbury". By some curious oversight, this latter remedy for monopoly powers fails to survive in modern legislation. The Statute of Anne provided that the term of protection was 14 years plus one final renewal of 14 years and it was used as a template for most other copyright legislation around the World including the US Copyright Act of 1790.

An incentive?

The Statute of Anne promoted the idea that most people think underlies modern copyright law: an incentive to authors to write books. Commercial exploitation however was the true purpose of the Act, as the Stationers promoted it and got their wishes.

The Copyright Act of 1814 raised the protection term to 28 years or the authors' life. Sergeant Talfourd's attempt to extend copyright to the life of the author and 60 years was frustrated by T.B. Macaulay in an inspirational speech in the House of Lords that repeats the themes in modern copyright [See the *Lesson From History* box over page].

Macaulay's *tour-de-force* single-handedly stopped the extension of copyright term until international treaties forced it 70 years later when the Copyright Act 1911 finally put paid to Macaulay's arguments by granting a term of protection of life plus 50 years. Around the turn of the 20th Century, Britain was a net exporter of copyright materials. As such, it had an interest in getting reciprocal protection for its authors and while this was possible with other European countries, the US was adopting a protectionist stance and granting protection only to its authors. The US is now a large net exporter and has lost the enthusiasm for protectionism since it now needs to protect multi-billion dollar exports.

Increasing copyright term was a global trend. The US copyright regime copied the English one: under the Copyright Act of 1790 the protection term was 14 years and one 14-year extension. Prior to that, copyright was perpetual and regulated by common law alone. The US term has changed forty times so far. More recently by the

1976 Copyright Act and most recently the Sonny Bono Act that extended copyright to life plus 70 years. For corporations it was raised to ninety-five years as a result of lobbying, together with substantial donations to politicians by the MPAA and Disney, which was concerned that some Mickey Mouse cartoons were about to pass into the public domain. The power to protect copyright is stated in the US Constitution as being: "... To promote the Progress of Science and the useful Arts...". There is no limit specified on the term and Congress feels it can decide whatever term it wishes. *Eldred v Ashcroft* [box, next page] was an attempt by Laurence Lessig to question this power. The case was recently completed before the US Supreme Court and resulted in a fairly unequivocal defeat 7-2 with Ginsberg delivering a majority verdict that failed to address many of the underlying questions. The UK has no written constitution so arguments based on an inherent limitation of power are broken.

The motivation at the core of copyright is that it is morally wrong to plagiarise, a theme consistent between common law countries and continental Europe. In the continental system this is expressed in the concept of moral rights, or *droit d'auteur*, that ascribe to the author the right to prevent mis-attribution, whether accidental or plagiaristic. However,

"When copyright is used abusively, it's hardly surprising that people question the value of intellectual property law."

practical limits exist on the power of law to enforce these moral constraints and at the same time preserve competition in a free market. Courts do not therefore extend a monopoly in the first embodiment of an idea; thus the author of *Midnight Commander* need not fear the authors of *Norton Commander* which inspired it. At a more pragmatic level though, a substantial part of the motivation was that to encourage the creation of new works, there often needed



ARTWORK BY DYLAN TEAGUE. JUDGE DREDD IS TM AND © REBELLION, 2003

Copyright



to be an entrepreneur to risk commercialising a work, encourage the author to produce more (and take a majority share of the profits). Everyone won; the author earned a living and a reputation, the entrepreneur got profit and the public got more works. Content producers raise the plagiarism argument frequently and governments and the public are generally responsive

“While plagiarism is abhorrent, it is an unavoidable fact of creative life. Indeed, sometimes the copy eclipses the original”



to it. While the notion of plagiarism is abhorrent, it is an unavoidable fact of creative life. Indeed, sometimes the copy eclipses the original. Ever heard of *Romeus and Juliet*? I doubt it, but it was this Arthur Brooke poem that inspired Shakespeare's *Romeo And Juliet*. Conan Doyle took *The Hound Of The Baskervilles* from a work of the almost unknown writer Fletcher Robinson (and is alleged to have murdered him when exposure was threatened). Disney (who now seek perpetual copyright) made their fortune by plagiarising the works of

the Brothers Grimm and other 19th Century authors. One of the expressed rationales for the GPL is to actively encourage copying. Richard Stallman argues that society would benefit from the ability to read the code of others then copy it, line by line if needed, or adapt it to provide a useful change. The success of Linux and the GPL is a telling justification of that idea.

The nature of copyright

So what is copyright? Broadly, it is a partial monopoly given for a period of time by the state to a person who reduces an original idea to some form of fixed expression. Thus, an original novel or computer program would *at the moment of fixation* give the author the right to prevent others copying a substantial part of it. By contrast, an impromptu speech would not attract protection since it would never have been fixed in writing. The main legal highlights of copyright then are originality, fixation and the type of work.

DURATION Copyright has the duration fixed by the statute applicable at the time the work was created. There have been several UK copyright statutes: 1911, 1956 and 1988 and their amendments.

Therefore to find the period term one has to consider the date of creation correlate it with statute and check retrospective modifications. Additionally, some types of work were unknown in previous legislation so one has to draw an analogy; for example films were treated like photographs for a while. Legislators now try to ‘future-proof’ statutory phrasing, thus storage media is not generally referred to as optical or magnetic lest someone make a device to store terabytes of data in an acoustic standing wave. The rules are quite detailed but for the most part it, in literary, dramatic or musical works the period is life plus 70 years.

FIXATION Fixation means making a non-ephemeral copy. Writing on paper is fixation, while writing on the surface of water is not. Fixation does not mean permanent: writing ones name in the snow would be fixation and writing ones name in the sky by plane would probably also be fixation. A more realistic example would be copying data into dynamic memory; even though the bit pattern may only last for a few nano-seconds courts regard it as fixation, so writing data into RAM constitutes fixation for the purposes of either creating or infringing copyright. This is something of

A Lesson From History

Reward authors for their literary labours



Macaulay's speech in the House of Lords in 1841 responding to Mr Sergeant Talfourd's bill attempting to increase the copyright term to 60 years. Single-handedly it got the bill rejected 45 votes to 38. Read it and see why:

“The question of copyright, Sir, like most questions of civil prudence, is neither black nor white, but grey. The system of copyright has great advantages and great disadvantages;

“... The advantages arising from a system of copyright are obvious. It is desirable that we should have a supply of good books; we cannot have such a supply unless men of letters are liberally remunerated; and the least objectionable way of remunerating them is by means of copyright. Such men must be remunerated for their literary labour. And there are only two ways in which they can be remunerated. One of those ways is patronage; the other is copyright.

“... Copyright is monopoly, and produces all the effects which the general voice of mankind attributes to monopoly. That monopoly makes things

dear is certainly a theory ... If, as my honourable and learned friend seems to think, the whole world is in the wrong on this point, if the real effect of monopoly is to make articles good and cheap, why does he stop short in his career of change? Why does he limit the operation of so salutary a principle to sixty years? Why does he consent to anything short of a perpetuity? He told us that in consenting to anything short of a perpetuity he was making a compromise between extreme right and expediency. But if his opinion about monopoly be correct, extreme right and expediency would coincide.

“I believe, Sir, that I may with safety take it for granted that the effect of monopoly generally is to make articles scarce, to make them dear, and to make them bad ... It is good that authors should be remunerated; and the least exceptionable way of remunerating them is by a monopoly. Yet monopoly is an evil. For the sake of the good we must submit to the evil; but the evil ought not to last a day longer than is necessary for the purpose of securing the good.

“... the evil effects of the monopoly are proportioned to the length of its duration. But the good effects for the sake of which we bear with the evil effects are by no means proportioned to the length of its duration. A monopoly of sixty years produces twice as much evil as a monopoly of thirty years, and thrice as much evil as a monopoly of twenty years. But it is by no means the fact that a posthumous monopoly of sixty years gives to an author thrice as much pleasure and thrice as strong a motive as a posthumous monopoly of twenty years. On the contrary, the difference is so small as to be hardly perceptible ... And a man is very little moved by the thought that in the year 2000 or 2100, somebody who claims through him will employ more shepherds than Prince Esterhazy.

“The principle of copyright is this. It is a tax on readers for the purpose of giving a bounty to writers. The tax is an exceedingly bad one; it is a tax on one of the most innocent and most salutary of human pleasures; ... I admit, however, the necessity of giving a bounty to

genius and learning. In order to give such a bounty, I willingly submit even to this severe and burdensome tax. Nay, I am ready to increase the tax, if it can be shown that by so doing I should proportionally increase the bounty. My complaint is, that my honourable and learned friend doubles, triples, quadruples, the tax, and makes scarcely any perceptible addition to the bounty.

“If, Sir, I wished to find a strong and perfect illustration of the effects which I anticipate from long copyright ... I should select the case of Milton's granddaughter. As often as this bill has been under discussion, the fate of Milton's granddaughter has been brought forward by the advocates of monopoly ... He has dilated on the sufferings, on the abject poverty, of this ill-fated woman ... But, Sir, will my honourable and learned friend tell me that this event ... was caused by the shortness of the term of copyright? Why, at that time, the duration of copyright was longer than even he, at present, proposes to make it. The monopoly lasted, not sixty years, but for ever.”

a legal fiction but without it copyright law would have problems with much of what goes on in PCs.

ORIGINALITY In its everyday meaning originality implies a quality judgment: novelty, imagination and great insight are called original. This approach is seen in Germany and various countries which regard copyright as a reward for creativity and skill. In common law countries (mainly the UK, US, Canada, South Africa, Australia, NZ and Ireland) the approach is pragmatic with lower barriers to obtaining protection. What is required in the UK for copyright protection is described by various phrases such as "sweat of the brow", "skill, labour and judgment" and displaying the entrepreneurial basis of copyright law: "labour skill and capital". All these mean no more than there has to be a certain minimum of effort in creating the work and that the work should not be copied. Cutting and pasting large chunks of run-of-the-mill text will often be insufficient; thus reproducing train timetables has been found insufficient. However compilations involving more effort can be protected. For example in the ever-changing tapestry of Red Hat's licence, it claims copyright for its recent distributions compilation. By taking a collection of RPMs, tarballs and such from many sources and burning them into an ISO they claim to have created a work. Whether they are right is a matter of conjecture; maybe they are not exercising enough skill, labour and judgment, and that the effort involved in dumping those files into a directory tree is *de minimis* and too insignificant. Maybe not. That said there is no question but that they have copyright in other elements of their distribution such as installers, utilities and their contributions to GNOME and KDE.

ORIGINALITY IN DATABASES

Databases often fall victim to the consequence of this attitude since, for example, a database of postal addresses is no more than a collocation of publicly available information. In the US case of *Feist v. Rural Telephone*, Feist copied Rural Telephone's database and was sued. However the court said there was no originality in the creation. The collection was merely one of facts that were not the substance of (or protected by) copyright law, as they were merely discovered and not

created. This caused great concern to the database industry, since the financial investment in creating a database can be huge. The response of the European Union to lobbying on these issues from the database industry was the creation of the Database Directive to protect the investment risk. That Directive created a *sui generis* (of its own kind) right, specifically for databases, that protects these otherwise unprotectable non-original databases. The definition of databases is remarkably and controversially broad: 'a collection of independent works, data or other materials that are arranged in a systematic or methodical way, and are individually accessible by electronic or other means'. There needs to have been a substantial investment in making the database, and this investment need not be just financial. The basic protection term is 15 years and up to 30 in certain circumstances. It is also possible to extend it perpetually if the database is substantially amended, such as by amendments, additions, deletions or similar. The database right is infringed by the substantial extraction or reutilisation of its contents, however repeated systematic extraction of insubstantial portions is treated as substantial extraction. Only databases of entities domiciled in the EU are protected. Various defences exist, such as fair dealing.

The Directive has been criticised as overreaching and unbalanced; for example, the criticism of fair dealing exceptions is that they are profoundly curtailed to protect the interests of database owners over the public. In addition, the length of protection is in practise perpetual – there's no encouragement of a public domain. The effect on scientific research both public and private would be badly affected, since research groups rely particularly on collations of public domain data. To magnify the issue yet further, the US wants to implement its own even tougher Database Directive, using as justification the EU Database Directive – *deja vu* anyone? At the heart of the Database Directive is the troubling proposition that it is the investing of money, and nothing else, in a venture that entitles entrepreneurs to a monopoly protection.

Eldred v Ashcroft

DCMA – unconstitutional?



In the mid 1990s, Eric Eldred conceived the idea of an online resource for providing high quality copies of public domain works. These he felt would improve on the often poor-quality resources then available. Based on a server at the distant end of a cable modem, he realised this objective and all was going well until the Sony Bono Act snatched back from the public domain those works that were available freely to all.

Laurence Lessig, a law professor at Stanford University, views this as a challenge to the core concept of copyright and he aims to wrest back for the public those works returned into copyright. He hopes also to mould future policy in the digital copyright arena as well and thereby to re-balance to policy objectives of modern US copyright law. Rejecting the trend of ever-extending copyright periods, 'perpetual copyright on the installment plan', he would set upon US entertainment corporations the obligation to return to the public domain their versions of that which they took from it in the first place. For instance, Disney took his inspiration from the fairy tales and folk lore of Europe and, presumably, Lessig regards the company as impertinent to keep it.

Rather than wait with thumbs a-twiddling for a suitable plaintiff, Lessig has sought one out and in Erik Eldred he succeeded. Eldred's daughters were set Nathaniel Hawthorne's *The Scarlet Letter* at school and they disliked it. Eldred sought to better present it by using new fonts, annotating the text with notes, hyperlinks, adding contemporary reviews and related works. He also tracked copyright expiry dates in other modern classics hoping to add them also. He planned to add

Sherwood Anderson's *Horses and Men* and *New Hampshire* by Robert Frost. These were printed in 1923 and due to expire in 1998 but the Copyright Term Extension Act stopped him. In his suit against the Attorney General, John Ashcroft said that the Act diminished the public domain, and he argued, in paraphrase, that;

1. It failed to encourage new works since its retrospective effect was not effective in encouraging a dead author to create new works.
2. By repeatedly extending copyright congress failed to draw a limit on the term of copyright as required by the US constitution.
3. It violated the First Amendment.
4. The act violated the Public Trust Doctrine by transferring public domain works into private hands without extracting a compensatory benefit.

The case was recently heard. Lessig was bullish about the prospects but as he self-deprecatingly admitted, he had lost all the other cases he had taken to court! Sadly his form was consistent and he lost: 7 judges ruled against him with a minority of 2 agreeing. The majority verdict was a bald and uncompromising assertion that Congress was entitled to do as it wished so long as the copyright term remained limited and the copyright landscape unchanged. The underlying assumptions of the Lessig challenge were barely addressed. That verdict rules out any hope of a further constitutional challenge, but while a sad and painful wound, all is not yet lost across the Pond. Hopes still flicker that resurrection of copyright may be unlawful, and there is the tempting prospect that the judgement of Ginsburg *et al* supports the view that the DMCA is unconstitutional.

In relation to software, the Database Directive has several ramifications: While software per se is excluded from the Directive, it is not clear that elements of software which can properly be regarded as a database can be also excluded. APIs probably fall within the definition of a database, as they are a methodical systematic collection of information and data. In the past, litigation has usually resulted in defeat for those companies trying to restrict use of access to them by copyright law. The Database Directive may therefore compromise the ability to reverse-engineer software and achieve



Copyright



interoperability, not withstanding other directives permit it. Concerns have been raised in the Debian community that wordlists may be subject to the new right and that non-EU packages be maintained – this may be a growing concern for distros in future. Further non-obvious consequences concern the freelance Unix sysadmin. Over a period of time, many collect bug and security patches to apply to systems they administer. These patch collections will fall within the definition of a database, and the freelancer will have a separate database right in the collection unrelated to any copyright issues.

Many contracts do not yet specify the contractual rights in these databases, since the issue is one that is 'under the radar' of many companies. Where this is so, the advantage lies with the freelancer. Unlike the employee, where all work is assumed to be the property of his employer, there is no matching assumption that works used or created in the course of an independent contractors engagement belong to those engaging him. That said, patches are usually applied only once rather than repeatedly, unlike many database records. Thus the practical ramifications are perhaps not great, at least until the freelancer leaves with his patch collection and the system needs reinstalling.

Infringement of copyright

Copyright gives a right to the owner of a copyright work to prevent reproductions or performances of the work. Ownership of the physical work is not the same as ownership of the copyright of the work embedded within it. Being the owner of a CD means ownership just of the physical disk, not the rights in the music or software on it.

COPYING It is for the copyright owner to prove that the work has been taken by the alleged infringer. Mere publication of an identical or similar work does not suffice. Why? Copyright is a partial, not a total, monopoly. Two people might independently create very similar works, but copyright merely prevents one taking another's work, not using one's own. By contrast, a patent generally prevents another using even their own creation. If an owner can demonstrate a credible nexus between his work and the appearance of the other work, the infringer may then have to disprove copying; for example, given the availability of source code to FOSS software, the sudden appearance of very similar code and functionality within a proprietary product may easily give rise to a presumption of copying. Microsoft for instance are well aware of this, and it is alleged that they go to some lengths to ensure that no one reads GPLed code before writing theirs. Others are less scrupulous, although many have been caught.

Copying can be subconscious, that is, where someone has been exposed to a work and then later reproduced it in the genuine belief it is their own. The court must decide if the similarity is coincidental and the degree of exposure to the original work. Similarly, indirect copying can occur by over-close imitation even without line-by-line copying. This is one of the most frequent allegations made in software infringement.

What constitutes copying depends to some extent on the nature of the work. Works can be artistic, literary, dramatic or musical; and what is infringing copying in one form of work is not necessarily so for another. For example, turning a 2D picture into a 3D object is an infringement of an artistic work, but turning a sound recording into picture might not be. The

rules are rather complex so I will not go into them but the CDPA 1988 specifies them [see *Resources List* opposite].

SUBSTANTIAL TAKING - THE IDEA/EXPRESSION DICHOTOMY

Taking a minuscule part of a work is not enough, what is taken must be substantial, *ie* above an undefined minimum. In lawyer-speak, it has to be more than mere *de minimis*.

Reproducing the above heading would probably be copying so small and insignificant that to attach liability to it for copyright infringement would be petty and unfair. There is a desire not to trivialise infringement and recognition that creativity would be stifled without creators having the ability to draw inspiration from others' ideas. This is reflected in the idea/expression doctrine. This says that at one end, no one can prevent others from the use of a general idea. At the other end, slavish word-for-word copying is the taking of an expression of that idea. Agatha Christie had no monopoly on the idea of writing books detailing the adventures of an individual investigating crime. However, her estate can prevent someone copying *Murder On The Orient Express*. In between these fixed posts lies a grey area in which the question is "which post is the defendant nearer?" This imprecision functions as a means by which the court can do justice to any claimant or defendant.

The judge merely declares, on the facts of the case, that what was done was/was not a taking of the expression. If enough is done to the original work, it may be said that no infringement occurs, because nothing sufficient of the original remains in the copied work. Equally, if what is copied was trivial, it may be said that exact copying is needed to infringe. The courts then assess the skill, labour and judgement of both defendant and plaintiff in setting the boundary.

LIABILITY FOR THE INFRINGEMENT OF OTHERS

The law will sometimes impose liability on someone other than the primary infringer. It is currently seen in high profile in P2P litigation. Usually the allegations are of vicarious and contributory infringement and in the UK 'authorising infringement'. Collectively these are known as secondary infringement. Typically, there is commercial dealing with a work where

Copyright extension

The public domain shrinks even further



The premise of Lessig's position in *Eldred v Ashcroft* was that the 1976 US Copyright Act and the Copyright Term Extension Act were exemplar patterns for all US copyright extension (and likely to form the basis of US pressure on the EU to extend term in later decades). What then was the origin of copyright extension? Regrettably the answer is Europe and the 19th century expansion of Lockean possessive individualism: that an author had market based rights that transcended the interests of society. Copyright was a property just like a house. A house doesn't fall into public possession after 28 years, so why should a work? This is the MPAA line. The US Constitution prevents indefinite copyright so they say copyright should last forever minus a day. Since infinity minus one is also infinity, it demonstrates that they should perhaps buy a few more mathematicians and fewer politicians.

The term was originally 14+14 years and Sergeant Talfourd's attempt in 1841 at extension were defeated by Macaulay. But, although defeated, he seeded the trend for extension which the Berlin revision of the Berne Convention in 1908 set at life plus 50 years. It was this the US felt obliged to synchronise within its 1976 Copyright Act by a twenty-year extension. Now things became ironic. In 1965 during the process of copyright reform, Germany planned a system of patronage for worthy talent. This was opposed as socialism and as a throwaway concession, an extension of twenty years to copyright term was granted. Then the term was raised to twenty years throughout the EU in the name of harmonisation with the Germans. The Bono Act was in turn justified as necessary to harmonise with the EU, conveniently forgetting no other country in the World nor any copyright treaties had this term.

the copies in question are known or reasonably believed to be infringing. In the UK this means importing, possessing in the course of business, selling, letting for hire, offering or exposing for sale or hire, exhibiting in the course of business and distributing in the course of business or otherwise distributing to an extent prejudicial to the copyright owner. Prior to the DMCA, the latter phrasing was not in American copyright law, so people were able to redistribute copyright works for free usually without liability.

Contributory infringement and vicarious liability is most often alleged when employees do some wrongful act and it is easier to go after the employer than the employee, since the employer will be richer. They arise in their most developed form in the US, where someone assists the primary infringer. The RIAA *et al* realised that these are an effective way of dealing with P2P. If one cannot easily attack the sharers then go after their employers, colleges, providers of bandwidth and the creators of sharing programs. Napster ultimately met its demise under these doctrines. The knowledge of the alleged secondary infringer is critical: what is 'knowing or having reason to believe'? Was it an active knowledge, a guilty belief? This was a fairly substantial burden for those suing or prosecuting and courts have weakened the burden of proof necessary by reading this as "notice of facts such as would suggest to a reasonable man that a breach of copyright was being committed". Put more colourfully, 'shutting ones eyes and ears to the obvious' perhaps?

Exceptions to copyright

Copyright is limited by exceptions to reflect public policy. The exceptions granted under so called 'fair dealing' provisions are much more limited than the similar sounding 'fair use' rights of US and Australian copyright law. These contain an unqualified test of 'fairness' rather than a definitive list of exceptions to copyright. UK copyright is thus much more owner-friendly. Nor are there the general private use provisions of continental European law.

The exceptions are too detailed to outline fully, but in brief, the exceptions exist for research or private study,

reporting current events, and criticism or review. Librarians, archivists and educational establishments get enhanced rights, reflecting the argument that since the public gives the exclusive right, it gets the benefit for public interests it values highly.

Additional to the statutory exceptions the courts retain the discretion to refuse to enforce copyright, where circumstances raise a judicial eyebrow. In the *Spycatcher* case, copyright in the memoirs of MI5 agent Peter Wright was said by a petulant-sounding House Of Lords to be unenforceable, so that none of the guilty could profit. Pornography might be in a similar position. Similarly, if there is a public interest in revealing hypocrisy or misdeeds, the use of copyright to prevent dissemination of that knowledge will not be permitted. The courts have said, for example, that where a celebrity presents a wholesome face to the world the tabloids may reveal facts like drug and alcohol abuse.

Changes in copyright law in the Internet era

The Internet has come out of nowhere for most of society. Originally just a military project that transformed into a channel for scientific institutions to collaborate, it began to mutate out-of-control from the point at which the public got access. The reaction of politicians and content producers was like that of animals on meeting something strange: "can I eat it, can it eat me or can I breed with it?". No. So there was no interest until popularity and compression technologies threatened vested interests. JPEGs enabled pornography exchange in various forums, which made politicians and the media excitable. Then MP3 and DivX compression enabled music and videos to be become shareable and content producers also became excitable. This new thing might just eat them. Content producers generally hate and distrust technology capable of copying – recall the campaigns aimed at stopping home audio and video taping. They reserve to the Internet a particular loathing based on the unique technical features that distinguish it from tape: it can instantly and at zero nominal cost create and distribute non-

degrading digital copies. The WIPO Copyright Directive was their response.

WIPO and copyright law

In the mid 1990s, there used to be a quaint and charmingly naive view among geeks that the Internet was essentially ungovernable and immune to effective censorship. This view was based on the resilient nature of packet routing, the inventiveness of geeks and the global nature of the connectivity. Indeed John Gilmore famously observed that "the Internet interprets censorship as damage and routes around it". The modern trend of Internet legislation began with WIPO, the UN body that negotiates international IP treaties, and is widely criticised for being in the pockets of American content lobbyists. True or not, the lobbyists devote considerable resources to influencing it. WIPO introduced the Copyright Treaty in 1996 that largely consisted of lobbyist representations, where the absence of consumer representation speaks volumes. The Treaty forms the basis for the underlying legal policies of signatories such as the UK, US, EU and many other countries. The most prominent features of the Treaty are in Articles 8, 10, 11 & 12 which deal, respectively, with the right to communicate works to the public,

“WIPO negotiates international IP treaties, and is widely criticised for being in the pockets of American lobbyists.”

exceptions, technological measures and rights management information.

In the European Union, the Treaty is enacted in Articles 3, 5, 6 and 7 of the European Union Copyright Directive (EUCD) which lifted much of the text word-for-word. Article 3 introduces an exclusive right to communicate a work to the public something UK copyright law has not seen before and its implications are unclear.

Article 5 provides a copyright exceptions clause that specifies nothing inconvenient to content holders and which fails to adequately meet issues of fair dealing.

Resources list



Eldred v Ashcroft

<http://eon.law.harvard.edu/openlaw/eldredvashcroft/>
<http://eldred.cc/legal/supremecourt.html>

Eldritch's site

<http://www.eldritchpress.org/>

Lessig's blog

<http://cyberlaw.stanford.edu/lessig/blog/>

Open law

<http://eon.law.harvard.edu/openlaw/>

Copyright activism

<http://ukcdr.org/> - campaign for digital rights

<http://www.geekivism.org/>

<http://www.fipr.org>

<http://www.creativecommons.org/>

The TCPA petition

<http://www.petitiononline.com/endtcpa1/petition.html>

European Union IP portal

<http://europa.eu.int/scadplus/leg/en/s06020.htm#AUTEUR>
 database directive text
<http://europa.eu.int/ISPO/infosoc/legreg/docs/969ec.html>

Copyright Designs and Patents Act 1988

http://www.legislation.hmso.gov.uk/acts/acts1988/Ukpga_19880048_en_1.htm

Privacy and the EUCD

<http://www.geeklawyer.org/EUCD-v-privacy.html>

Copyright



Content creators have long put faith in technical means of preventing copying. Their maxim is 'the answer to the machine is in the machine.' This has led to circumvention, since many find the security inconvenient and intrusive. Legislators have stepped in to prevent this; in many countries, it is already illegal to remove copy protection capabilities. In the UK, section 296 of the Copyright Designs and Patents Act 1988 provides similar protection, but it will be substantially amended and extended to comply with the EUCD. Article 6 bears on this by requiring states to protect effective technical measures. Circumvention hardware and software such as *deCSS*, Xbox modchips and DVD players adapted by a third party to be omni-regional probably fall within this. It's worth noting that the meaning of the word 'effective' is not likely to be taken to refer to the quality of the protection: ROT13 encryption will be enough. The courts are going to say that as long as it achieves, or attempts to achieve, the intended effect it is 'effective' and that protection quality is immaterial.

Anti-competitive

The total inability to circumvent technological protection measures (TPM) for legitimate purposes such as fair dealing is controversial. The EU was well aware that TPM could create a perpetual copyright, since if the system is sufficiently effective and uncrackable, the only way to obtain access would be a licence on whatever terms they could extract from the hostage consumer. Similarly, TPM are seemingly being used for anti-competitive purposes. It has already been alleged that some mobile phones are programmed to flatten non-manufacturer batteries fast if they fail a challenge-response protocol. At present, cloners can probably legally reverse-engineer the protocols. However, corporations have twigged that they need only wrap fatuous copyright protection into this loop and cloning will be outlawed. How courts will react is too hard to predict. There is a credible but uncertain anti-competitive behaviour argument under Articles 85 or 86 of the Treaty Of Rome. It would have been much better to write it into the Directive. I and other members of the Campaign for Digital

Rights have made this point in meetings with the Copyright Directorate at the UK Patent Office. How they will react will be seen in impending UK copyright law amendments.

Article 7 protects rights management information (RMI) which would be things like digital watermarks. Once one begins to encapsulate digital works within a protection framework, it becomes very desirable to know who is using your works in what manner, where and when. A long-felt desire of content producers is to have this information. It would enable them to protect their works, but more importantly, it would have the potential to enable them to profile the user. Once this can be done, no longer would they be compelled to indulge in the economic averaging process that currently leads to one fixed price for all work, such as £12 for a CD, but rather a price dependent on what job you have and your net monthly income. In this dystopian vision, not only the quarterly electricity bill would arrive; so also the, as yet mythical, Copyright Users Society bill that would list who in your family used what and what the pricing is. The means for doing this will be rights management systems, that will 'phone home' via the Internet. When they do, they may also disclose personal information. While the directive mentions this and requires privacy protection, it does not, unlike the DMCA, explicitly permit you the right of self-defence in the face of violating RMI systems, and it is unclear if circumventing a standalone RMI system is lawful. If it is you can be assured they will be coupled tightly with the TPM. Without lobbying and protest, this will become the era of the digital glass cage, of copyright spyware where your music watches you. This is one of the more controversial parts of the EUCD (see *Resources List* for more).


The future

Geeks bless the coming of the Internet as an era where communication and freedom built a community free to exchange code, ideas and values. What a cruel irony if that very network seeds the destruction of those values through corporate and government paranoia and lack of insight. Much of this is creeping up on the public unbeknown to them. Magazines give away free

trojan music that requires upgrading to the latest *Windows Media Player* with its DRM spy-ware: where a supposedly anonymous GUID sends back information to servers at Microsoft and the RIAA to determine if you can listen to 'your' music.

What are the implications of this? It is entirely likely that we will see the same issues arising here as in the US: laws targeted at copyright infringers will instead be used to silence those with forbidden knowledge, maybe a UK Sklyarov case with a security researcher being threatened by an insecure copyright owner. Perhaps with Ross Anderson in the role of Edward Felten. The knock-on effect is likely to be more serious, since security research, despite supposed protection for it in the EUCD, will be prevented or curtailed 'just to be on the safe side'; good for crackers and GCHQ, but no-one else. It may also change the business model for music: pay on demand, with privacy compromised to satisfy the control urges of the content industries. Perhaps Linux and all other FOSS software will be subject to attempts at curtailment, to protect the security of copyright protection subsystems. By pursuing this inexorable logic of legal and technical control, content industries signal to customers what it thinks of them: that they are thieves to be patted down by bouncers lest they leave the Internet with digital contraband. And secondly, that they are a mere resource to be milked frequently.

The copyright fight is part of a larger fight to control the future of digital systems – a wider issue than music copying involving also information control, censorship and accessibility. It threatens an era when we exchange democratic control of copyright for corporate control implemented by technology.

Copyright laws benefit all of us, providing that they rigidly enforce the balance of rights and interests between consumer and producer. Governments, through either corruption, conservative thinking or an unwillingness to say no to corporations, are letting this valuable right slip inexorably into disrespect by permitting its objectives to be distorted by the self-interested. It need not and should not be like that, and it is up to geeks to fight it on behalf of those not aware battle has been joined. 

About the author

David Harris is an IP barrister who idles away his days at www.ukitlaw.com waiting desperately for a clued-up client to ask him something about Linux rather than software patents. When not suing people, he writes 'depressingly poor code'.

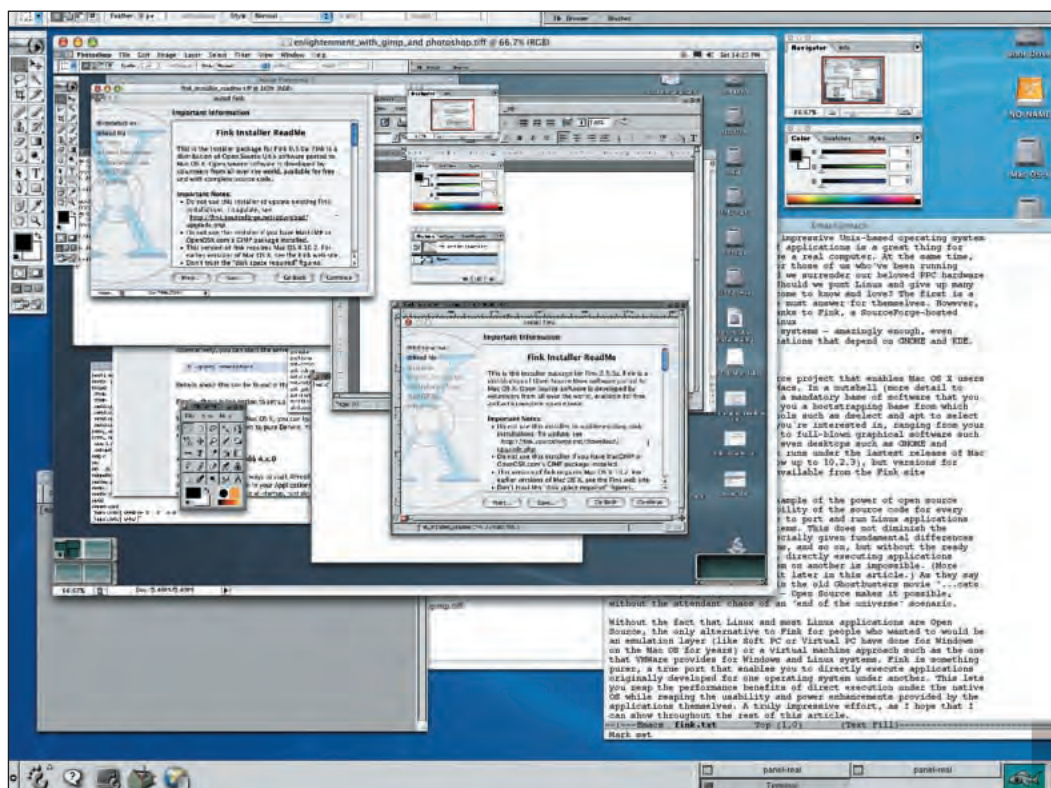


Bringing Linux to Mac OS X

A BITE OF THE APPLE

Bill von Hagen shows that Linux will work on just about any computer system, even those that favour a mouse with only one button.

Photoshop and GNOME – working side by side on the Mac.



The emergence of Mac OS X as the official Macintosh operating system from Apple – an impressive Unix-based operating system with a real GUI and thousands of applications – is a great thing for both Macintosh users and Unix lovers everywhere. At the same time, it poses a conceptual problem for those of us who've been running Linux on Apple hardware – should we surrender our beloved PPC hardware to a commercial OS and vendor? Should we punt Linux and give up many of the applications that we've come to know and love? The

first is a philosophical question that I wouldn't touch with a branding iron. However, the second is largely a NOOP thanks to *Fink*, a SourceForge-hosted project to make "traditional" Linux software available for Mac OS X systems – amazingly enough, even including graphical Linux applications that depend on GNOME and KDE.

What is Fink?

The *Fink* project is an Open Source project that enables Mac OS X users to run Linux software on their Macs. In a nutshell (more detail to follow), the *Fink* folks provide a mandatory base of

software that you install on your Mac. This gives you a bootstrapping base from which you use standard Debian Linux tools such as `dselect` and `apt` to select and install the Linux packages you're interested in, ranging from your favourite command-line utilities to full-blown graphical software such as *XFree86*, window managers, and even desktops such as GNOME and KDE. The current version of Fink runs under the latest release of Mac OS X (Jaguar – Mac OS X 10.2, now up to 10.2.3), but versions for older release Mac OS X are available from the *Fink* site (<http://fink.sourceforge.net>).

The *Fink* project is a stellar example of the power of Open Source software – thanks to the availability of the source code for every Linux application, it's possible to port and run Linux applications natively on other operating systems. This does not diminish the wizardry required to do so, especially given fundamental differences in graphics libraries, subsystems, and so on. However, without easy (and owner-authorised) access to the source code, directly executing applications designed for one operating system on another is impossible. (More about Open Source and Apple a bit later.) As they say regarding the end of the world in the *Ghostbusters* movie "...cats and dogs sleeping together?..." – Open Source makes it possible, and without all those rude bits from Revelations or the Apocalypse.

Without the fact that Linux and most Linux applications are Open Source, the only alternative to Fink for people who wanted to would be an emulation layer (like Soft PC or Virtual PC have done for Windows on the Mac OS for years) or a virtual machine approach such as the one that VMWare provides for Windows and Linux systems. Fink is something new, a true port that enables you to directly execute applications originally developed for one operating system under another. This lets you reap the performance benefits of direct execution under the native OS while keeping the usability and power enhancements provided by the applications themselves. A truly impressive effort, as I hope that I can show throughout the rest of this article.

“Fink is a stellar example of Open Source — thanks to the availability of source code, it’s possible to port and run Linux apps natively on other operating systems.”

Apple and Open Source

APSL isn’t necessarily Free Software

Since the inception of Mac OS X, Apple has made a great deal of noise about the fact that significant portions of OS X are Open Source, which is perhaps a little misleading. It actually means that they are available under the Apple Public Source License (APSL). The APSL can be viewed at the URL

www.opensource.apple.com/apsl.

Rather than boring you with them here, you can get detailed information about the APSL and its implications at the URL <http://developer.apple.com/darwin/ps-faq.html>. To minimise confusion, we’ll refer to both APSL and Open Source applications as Open Source throughout

the remainder of this article, though there are important legal distinctions.

The portions of Mac OS X that are Open Source are largely confined to the Mac OS X kernel, known as *Darwin*, and a variety of kernel-level libraries and interfaces. What most users think of as Mac OS X, the GUI (known as *Aqua*) and the Mac OS 9 emulation layer, known as *Carbon*, are not Open Source. Frankly, these are the keys to the Apple kingdom, and Apple is justified in keeping them as its intellectual property. Without these, Apple is just an insightful vendor of sexy, powerful, and (some would say) over-priced hardware.

that VMWare provides for Windows and Linux systems. *Fink* is something purer, a true port that enables you to directly execute applications originally developed for one operating system under another. This lets you reap the performance benefits of direct execution under the native OS while reaping the usability and power enhancements provided by the applications themselves. A truly impressive effort, as I hope that I can show in the rest of this article.

Why is this article appearing in a Linux magazine? Because, as Richard Stallman will be happy to see in print again, what most people think of as Linux is far more than just a whizzy, Open Source kernel. To a desktop computer user, an operating system by itself is essentially useless — it’s the applications and utilities that surround an operating system which make it possible to actually get work done (or indeed play, for that matter). This is the core of the reason why most people who are willing to pronounce the extra syllable refer to Linux as

“GNU/Linux”. Much of what we now know as Linux existed way before Linus Torvalds. Without such GNU applications such as bash, emacs, gcc, and enhanced versions of every traditional Unix utility, other Open Source applications like *XFree86*, GNOME, and KDE, and the thousands of graphical applications that depend on the last three, the Linux kernel would be little more than a clever, high-powered way to use up disk space on desktop computer systems.

Looked at in this light, the *Fink* project is not only a great solution for Linux lovers who also want (or need) to use commercial Apple Mac OS X applications, but it’s a great promoter and tremendous example of Open Source and the powerful applications that are available under Linux.

The remainder of this article provides some insights into how Fink works, explores the basics of installing and using Fink on OS X, and provides some truly amazing screen shots of well-known, graphical Linux apps running on a Mac OS X desktop.

How Does It Work?

Earlier articles in *Linux Format* (LXF 18) explored the basics of Mac OS X and provided comparisons with Linux in terms of its operating system, Unix roots, and the range of Open Source utilities included with the world’s most popular commercial version of Unix.

If you’re an OS X user, an initial concern about using Fink to run OS X and Linux software side-by-side is the problem of namespace collisions. What happens when you accidentally (or intentionally) install a Linux version of an application that is already distributed with Mac OS X? Luckily, the answer is simple and pleasant — nothing except using up a bit of additional disk space.

The key to almost all *Fink* project software is that it installs into its own directory hierarchy. This hierarchy is



TheFinkProject



rooted at the /sw directory on your system's boot disk. The only exception to this rule is the X Window system provided by *Fink*, which is the standard XDarwin (<http://www.xdarwin.org>) that is also freely available and independent of *Fink*. All modern X Window system distributions are rooted at the standard /usr/X11R6 directory because this path is hardwired into too many X applications, configuration files, resources, and so on.

Fink's use of its own directory hierarchy means that Linux

distribution, basing its installer on the impressive APT (A Package Tool) tools, the classic Debian dselect application, and the basic power of the Debian packaging system, *dpkg*.

Installing *Fink* is a simple bootstrapping operation. First, you download a standard Mac OS X installer for the core *Fink* applications and execute it the same way that you install any Mac OS X software. Next, you use the utilities installed by the *Fink* core to select and download whatever other *Fink*-based applications you want.

in your home directory:

```
source /sw/bin/init.csh
```


If you have reconfigured your OS X login so that you use /bin/bash or /bin/sh as your shell, you would add the following line to the .bashrc file in your home directory:

```
./sw/bin/init.sh
```

Next, open a new OS X Terminal window and execute the "fink scanpackages" application at the command prompt. This command does the last bit of *Fink* initialisation by installing a few packages and verifying the list of *Fink* packages that are now available on your OS X system. You're now ready to start using *Fink*!

NOTE: The *Fink* installer installs the minimum set of applications and utilities required to get started with *Fink*. Initially, *Fink* is entirely command-line oriented. You will only be able to execute graphical Linux applications under *Fink* on your OS X system after installing prerequisites such as X11R6 and the applications (and related libraries) themselves.

"Once you've installed and initialised Fink, installing Fink packages is as easy as installing apps on a Debian Linux system."



applications from *Fink* can happily coexist with any native Mac OS X versions of those same applications. Executing a specific version of an application is simply a matter of setting your path appropriately or executing such applications by their full pathnames. No fuss, no muss, as many Americans might say.

Installing Fink

As you might expect and hope, installing *Fink* is an incremental process. Given that there are thousands of Linux applications, installing something labelled "everything" would suck disk space even faster than Microsoft Windows. The *Fink* project takes advantage of the power of the Debian Linux

To install the *Fink* core on your OS X system, download the appropriate installer from the *Fink* download page <http://fink.sourceforge.net/download/index.php>. If you're running Mac OS X 10.2, also known as Jaguar (and you should be), the version of the *Fink* installer at the time this article was written was 0.5.0a. If you're still running Mac OS X 10.1, you should download the *Fink* 0.4.1 installer. These installers aren't trivial, each is an OS X-mountable disk image of 11 MB or so. These images each contain a Mac OS X

Once you've downloaded the appropriate *Fink* installer to your OS X system, double-click on the downloaded image to validate its contents and mount it as a virtual drive on your OS X desktop. You then open this disk and double click on the 'Fink Installer.pkg' icon. This displays a standard OS X installer as shown on the left. The remainder of the install process is the same as installing any other OS X application, except that *Fink* must be installed on the disk from which you boot Mac OS X.

Once you've completed the installation process, you only need to modify your .cshrc so that the correct environment variables are set for you so that you can use *Fink* with OS X. If you are using /bin/tcsh as your login shell under OS X (the default), you add the following line to the .cshrc file

Installing Additional Packages Under Fink

Once you've installed and initialised *Fink* as explained in the previous section, installing *fink* packages is as easy as installing applications on a Debian Linux system. *Fink*, like modern Debian Linux, uses Debian's APT (A Package Tool) interface to the underlying Debian package system. *Fink* provides two standard Debian utilities to simplify package retrieval and installation. These are the purely command-line *apt-get* utility, and the terminal-oriented *dselect* utility, which uses the classic curses/ncurses library to provide a user interface of sorts. (Its name stands for 'Debian Select' rather than the actual task that it unfortunately sounds very like.)

If you've never installed Debian Linux or any Debian-based Linux distribution, APT and related applications such as *dselect* are what Red Hat's RPM tools would like to be when they grow up. *Dselect* uses a special APT method to not only understand package dependencies, but retrieve relevant packages in the right order, install them in that sequence, and basically just "do the right thing" in Unix parlance. This section focuses on using the *dselect*

The *Fink* Core Installer for Mac OS X



utility, since this is the easiest *Fink* mechanism for retrieving and installing sets of packages. If you're a command-line fan and know what you're looking for, you can always use the *apt-get* utility to install additional packages. An impressive alternative to the command-line or terminal-oriented approaches offered by *apt-get* and *dselect* is the highly capable *FinkCommander*, a graphical, Mac OS X interface for the *Fink* packaging system. For more information about *FinkCommander*, see the box to the right of this column.

To install additional sets of *Fink* packages using *dselect*, execute the **sudo dselect** command from an OS X Terminal window. *Dselect* enables you to select packages for installation, see exactly what packages are already installed on your system, and so on. You can use the OS X arrow keys or the standard emacs ^p (previous) and ^n (next) commands to move up and down between various *dselect* options, and can also press the initial letter or associated number of each *dselect* command to select that menu. You then press **Return** to execute the specified command or enter a selected submenu.

The first thing that you may want to do when using *dselect* is to verify the default configuration provided by *Fink*. This tells fink the sites from which to obtain packages, etc. You can access this using the **A** command.

Next, it's time to peruse the *Fink* software archives, see what packages are available, and select those that you want to install to run on your Apple Mac's OS X system.

During installation (depending on the packages you've selected to

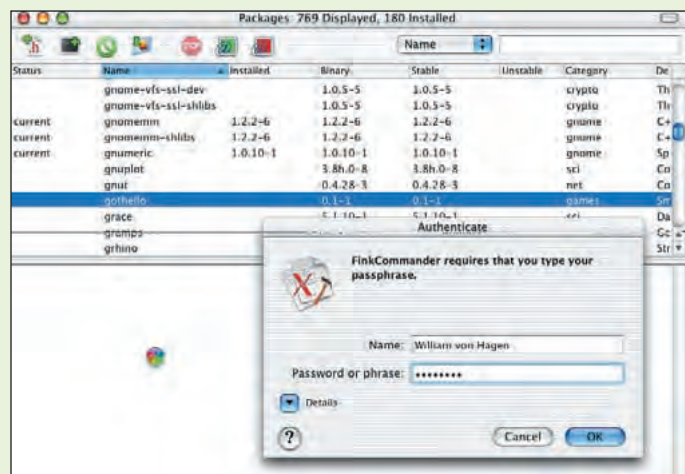
FinkCommander

A GUI for the Fink Package Manager

FinkCommander (<http://finkcommander.sourceforge.net>) is a graphical user interface for the *Fink* software packaging system that runs under Mac OS X. It provides an intuitive front-end to the *Fink* command-line tools for downloading and installing *Fink* packages – and is, of course, GPL itself.

The most recent versions of *FinkCommander* use Apple's Security Framework to obtain the authorisation necessary to run *Fink* commands. The first time you run *FinkCommander*, it prompts you for an administrative password for your Mac OS X system, and then configures itself to run as this user in the future.

Compared to the questionable elegance of *dselect*'s curses-based UI, *FinkCommander* is a simple but eminently useful utility. *FinkCommander* makes it easy to scroll through the list of available *Fink* packages, get information



FinkCommander, the Mac OS X Tool for Maintaining *Fink* Packages

about the status of those packages, see what's already installed on your system, and install other packages. Unless you're

a pure command-line aficionado (in which case – why do you have a Mac?!), *FinkCommander* is highly recommended.

install), you may see messages about suggested or mandatory modifications to the OS X distributed administrative database. Here's an example of some of these messages – the following user entries will be added to your NetInfo database:

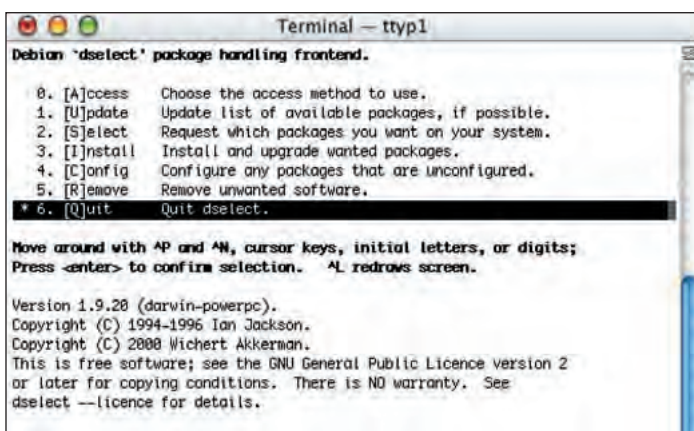
```
news:*:250:250::0:0:News Server:/dev/null:/dev/null
pgsql:*:252:252::0:0:PostgreSQL Database Server:/dev/null:/dev/null
games:*:253:253::0:0:Game Files Owner:/dev/null:/dev/null
canna:*:254:254::0:0:Canna Japanese Input Server:/dev/null:/dev/null
tomcat:*:257:257::0:0:Tomcat Servlet Engine:/sw/var/tomcat4:/dev/null
jabber:*:258:258::0:0:Jabber
```

Messaging Server:/dev/null:/dev/null

Messages of this sort mean that any existing NetInfo entries with these names, user IDs (UIDs), or group IDs (GIDs) will be overwritten when you install *Fink*. Some *Fink* packages will not work correctly unless appropriate user and group entries are present in the NetInfo database. If you're using a standalone OS X system and are not running services such as netnews, PostgreSQL, and so on, you can usually answer "yes" here. If it's a standalone machine and you don't know whether you're running these services, it probably isn't. If your system is running such services and they depend on specific UIDs or GIDs, you should check your NetInfo database carefully or consult your local sysadmin before arbitrarily updating the NetInfo database.

Installing Binaries or Building Your Own

The *apt-get* and *dselect* approach explained in the the previous section installs pre-compiled binaries of your favourite Linux applications on your Mac OS X. One of the most important implications of the Open Source movement is that the applications you may depend on are continually improving, and the enhanced source code is always available. Installing Linux applications on your OS X



Fink/Debian's dselect Utility.

TheFinkProject



systems using *Fink* therefore means that you're installing whatever binary snapshot was available when you installed, which isn't necessarily the greatest and is almost certainly not the latest version of the application.

The *Fink* project provides an alternative to installing binaries using Fink's *apt-get* and *dselect* tools, through a utility that actually obtains the source code for the applications you're interested in and builds them

on your system. This aptly named application is called "fink". The fink application runs from the command-line and uses parallel syntax to the most important command-line options for *apt-get*. For example, instead of running a command such as **apt-get install foo** to retrieve and install the package *foo*, you instead execute **fink install foo**, which retrieves the source code for the package *foo*, builds it on your system, and installs it for you.

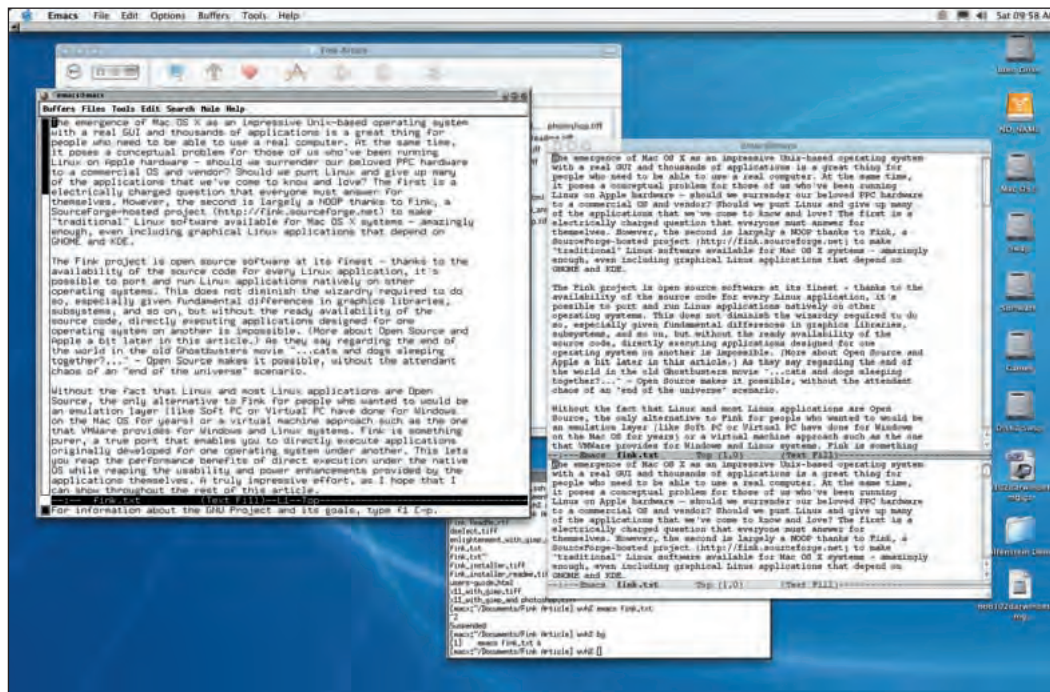
Using the fink utility, rolling your own Linux binaries for use under *Fink* on your OS X system is as easy as it could possible be. This may be especially significant to you if you're using newer peripherals on your OS X system and want the latest and greatest Linux-side support for those devices. This may also be especially important to you if you're using *Fink* with X11 implementations other than Fink's own, as will be explained in the next section.

X Window on OS X

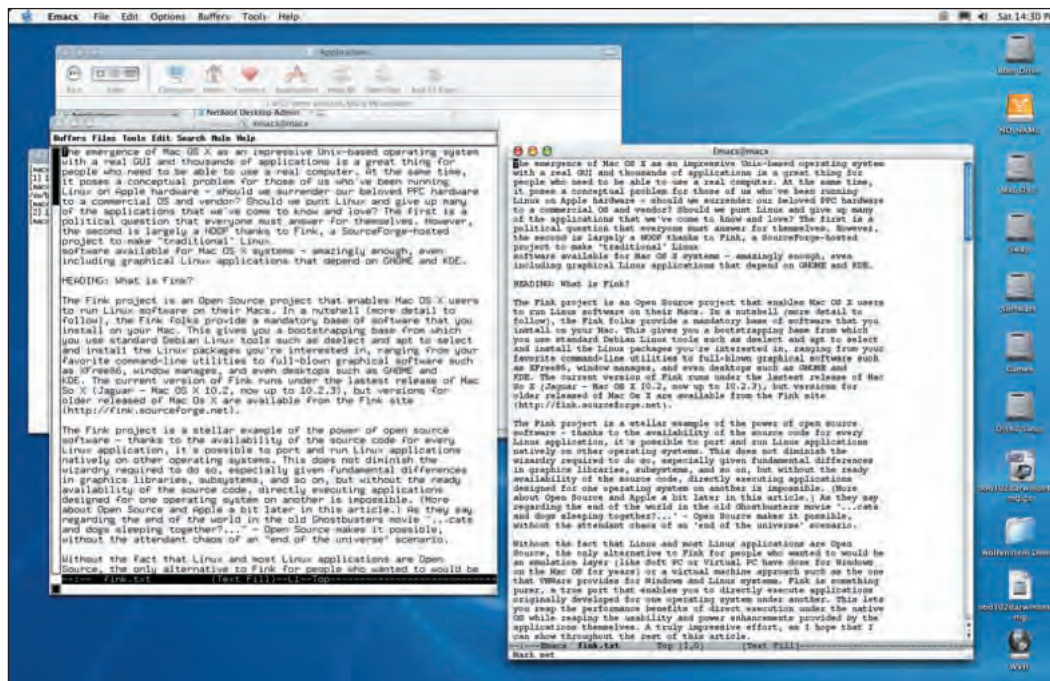
Support for the X Window system on Mac OS X is built into recent versions of *XFree86*. The X Window system server for Mac OS X is known as *XDarwin*, referring to the *Darwin* kernel used by Mac OS X. The development of *XFree86* for Mac OS X is done by the XonX project (<http://sourceforge.net/projects/xonx>), while the home page for the XDarwin X server itself and much of the end-user information about it can be found at the XDarwin home page (www.xdarwin.org).

Being able to run X11 on a system that already depends upon another GUI is an impressive feat, but X11 and OS X's *Quartz* graphics subsystem and *Aqua* interface coexist quite happily. XDarwin supports both a rooted mode, where all X Window system applications displays within their own virtual screen, and a rootless mode, where X Window system applications cohabit the OS X desktop, can be independently minimised to the dock, and so on.

XDarwin does not depend upon *Fink* – it's a totally separate project. You can download XDarwin from the XonX or XDarwin sites and run whatever the X Window system applications included with the base distribution without ever needing *Fink*. However, installing *XFree86* through Fink provides a number of advantages, primarily easy updating whenever new versions are released, and access to thousands of other X Window system applications. These range from other window managers (only *twm* is part of the base XDarwin distribution), GNOME, KDE, and all of their related applications. GNOME support is a tad more up-to-date than KDE support, but you still can't beat the price.



Emacs – Native OS X and Under XDarwin.

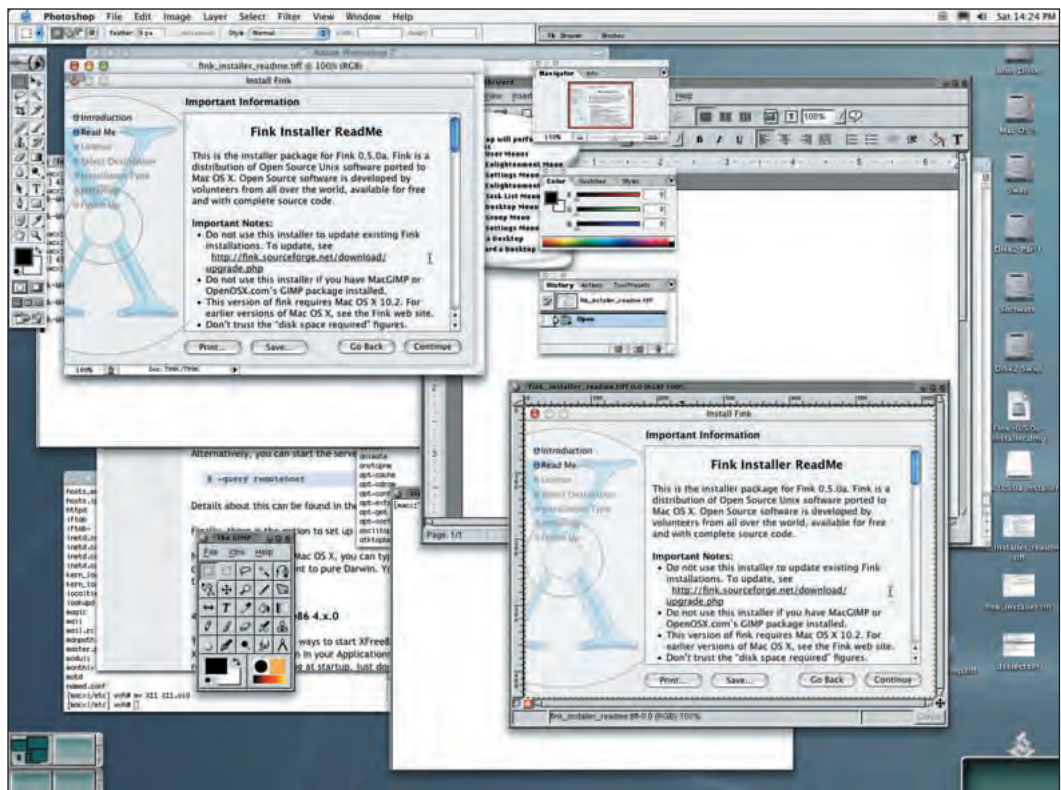


Emacs – Native OS X and Under Apple's X11.


If you've installed *XFree86* as a part of Fink, you can start the XDarwin X server from an OS X terminal command-line, using a command like **startx -- -rootless &** to start the X server in rootless mode, or **startx -- -fullscreen &** to start it with its own virtual screen. Installing XDarwin also creates a standard OS X application framework (XDarwin) in your /Applications folder. You can then start XDarwin by clicking on its icon like any other OS X application.

You may encounter a few glitches when using XDarwin in rootless mode, but these are only visual irritations, not functional problems. The most visible of which is the lack of visual feedback when you resize X11 application windows. Windows still resize correctly, but you don't get the rubberbanding or opaque resizing feedback that you would traditionally expect.

Very recently, on January 7, 2003, Apple released a custom X11 implementation based on *XFree86*



“Being able to run apps like The GIMP and Photoshop side-by-side makes it easier for users to migrate to Open Source.”



Fink, GIMP, and Photoshop – all in a day's work!

which includes *Quartz* rendering and accelerated OpenGL. This made the OS X user community very happy but probably was less than thrilling for existing vendors of X Window system implementations for Mac OS X such as Tenon Intersystems or eXodus. Statements on the Apple Web site say that Apple hasn't decided what to charge for the official release, but for now the beta is free. You can download it for yourself from www.apple.com/macosx/x11. When downloading this, you should also download and install the SDK listed at the bottom of the page. Apple's Mac OS X SDK enables developers to build most X11 R6.6 applications by simply recompiling under OS X.

Apple's X11 for Mac OS X uses Mac OS X's own *Quartz* graphics system, which provides hardware-accelerated 2D and 3D graphics. This provides faster text scrolling, dynamic

dragging and resizing of windows, and 3D animation thanks to OpenGL. Apple's X11 also comes with a special Apple window manager called *quartz-wm*, which is completely integrated with the *Quartz* and *Aqua* environments. If you're already running XDarwin, you should rename your *.xinitrc* file so that you can experiment with this window manager rather than the one you were using with XDarwin.

If you have existing *XFree86* packages installed, through *Fink* or otherwise, you can replace *Fink's* XDarwin with the Apple X11 server quite easily. Make sure that you first remove your old X11 packages, then install Apple's X11 and X11 SDK, and then run either **sudo apt-get update; sudo apt-get install system-xfree86** (for binaries) or **fink selfupdate-cvs; fink install system-xfree86** (to get and build the latest system-xfree86 package from CVS).

Detailed instructions on using the X Window system on *Darwin* and Mac OS X are available at <http://fink.sourceforge.net/doc/x11/index.php>.

Wrapping Up

Just like Linux and the various "BSD" operating systems, Apple's Mac OS X finally puts a Un*x box on your desktop. Philosophy and politics aside, there are work environments where Apple hardware and Mac OS X are mandatory. The *Fink* Project is an impressive piece of work that enables you to quickly and easily set up your favourite GNU/Linux applications and continue to use them while still getting work done in commercial Mac OS X applications. You may find that being able to run applications such as *The GIMP* and *Photoshop* side-by-side makes it easier for you to migrate to Open Source solutions to solve day-to-day work problems. [LXF](http://www.linuxformat.co.uk)

Some Useful Fink Links

If you *Fink* you want more...

Fink Project:

<http://fink.sourceforge.net>

Fink Installer for Mac OS X 10.2

(Jaguar):

<http://prdownloads.sourceforge.net/fink/Fink-0.5.0a-Installer.dmg?download>

Fink Installer for Mac OS X 10.1:

<http://prdownloads.sourceforge.net/fink/Fink-0.4.1-installer.dmg?download>

Fink FAQ:

<http://fink.sourceforge.net/faq/index.php>

Fink X11 FAQ:

<http://fink.sourceforge.net/doc/x11/index.php>

Where can I turn?



The Linux Newbie Administrator Guide aims to take you from decision to installation and beyond.



Linux has gained a reputation for being obtuse and support characterised as consisting of little more than a patronising “RTFM”. Well, it just isn’t so. There’s so much advice out there: **Andy Channelle** shows us where to look...

A computer operating system is a complex beast and no one can be expected to know about everything that’s really going on under the bonnet, so eventually, you’re probably going to need a little hand-holding. But where can you go?

If you’ve bought a boxed distribution from the likes of SuSE, Red Hat or Mandrake, your first port of call should be the manual. These are often graded to provide easy access to, for instance, a Quick Start guide, a more detailed system manual and, in the case of SuSE, a book covering a good selection of the apps that feature in the distro. The real problem with the printed word is that it ‘captures a moment’ that is soon passed and applications, desktops and OSs can change rapidly, which is where e-docs come in.

It has become common, in these cost-conscious times, to forgo the expense of a printed manual in favour of electronic documentation; and if you’ve downloaded an ISO or application from the Net, the chance of getting hardcopy documentation is

vanishingly small, unless you print it yourself from the readme files or help files that often accompany an app. Readmes tend to cover late-breaking issues, known bugs etc. and are an important resource to investigate before submitting a report on a bug you’ve discovered. The other options are to hit the Help drop-down that usually appears on the right side of a menu bar (if you’re using an app with a GUI) or combing through man pages.

The Help files are self-explanatory, but the **man** command goes back to the origins of Unix and is used – via the command line or something like **xman** which adds a GUI – to find, view and print online manuals. Typing, for instance, **man mount** at the command prompt will search for documents containing references to the ‘mount’ command. As with most search engines, you can limit your query to a particular section or keyword. Type **man man** for a list of all options for the command, as well as instruction on how to hone your searches and avoid the scatter-gun approach. Man pages are not written to win literary awards, they get straight to the point and are best suited to uncovering the various options/syntax of a command. They can also be used for pairing up a function with its associated command; just search using the function as a keyword.

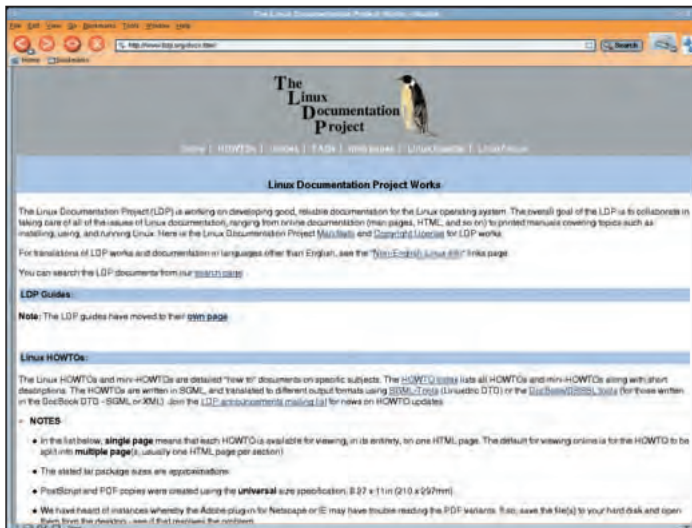
If you’ve exhausted the paper and online manuals, the next step in the quest for satisfaction should be the website of your distribution vendor. Boxed products tend to include some form of telephone or online support, with the most intensive options covering the installation and bedding down period (the first 60–90 days). Following this period, all of the big vendors offer a combination of free and

paid support varying from general configuration queries to highly tailored support designed for your particular network, hardware and software set up. The latter is obviously designed with enterprise and small business users in mind, but if you’re really stuck, a small ‘per-incident’ charge may pay dividends if it gets your system off the floor.

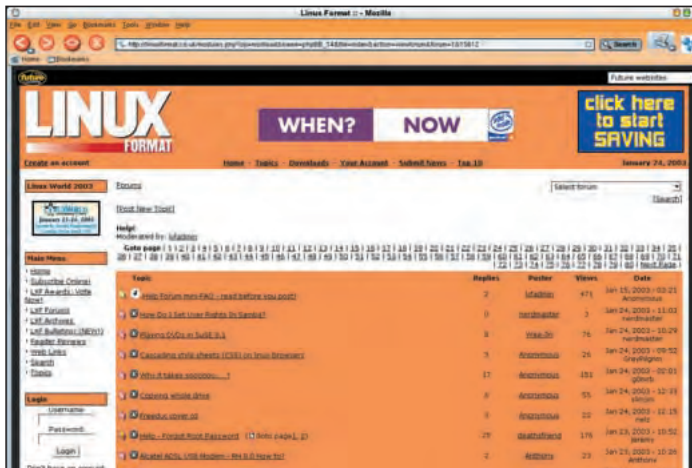
On your own?

Most vendors will also have some form of forum support system where other users and company experts may offer solutions to your problems. Though forum-based support is usually free and open to anyone, it does have its downsides: your answer may be a long time coming, for instance, and depend on a fellow user encountering in the past the same problem you have. Any potential solutions may also be figured out on hardware and software far different from what you’re using. In many cases you’ll get a timely response, but if your job, company or thesis depends on an answer, you may find paying for support is your best option.

For some users, solving a problem is part of the enjoyment (within reason) of using a computer. If this is you, browsing the Linux Documentation Project at www.tldp.org/docs.html will be a real eye-opener, as it contains not just a fully-featured set of man pages (with links to a couple of decent search engines), but also a vast store of Guides, Frequently Asked Questions and HOWTOs. The guides alone are a worthy addition to your support armoury, and contain information on everything from Advanced Bash Scripting to Kernel Module Programming.



The Linux Documentation Project is a central repository for hundreds of helpful documents, for everyone from newbies to experienced sysadmin.



The LXF forums are manned by knowledgeable and friendly Linux users.

Some of these guides are written with new users in mind, so don't launch into reams of coding without preamble. Of particular note is the Linux Newbie Administrator Guide <http://linux-newbie.sunsite.dk/>. It starts with very simple questions covering why you may want to adopt Linux in the first place, moves through acquiring a distribution, installation and boot issues, before covering the advantages of Open Source and proprietary software, networking and finally a user-friendly guide to kernel upgrading. If you're considering migrating to Linux, this is the next best thing to a very thick 'How To Do Everything In Linux'-style book. And of course, it saves you lots of money.

In between the terse man pages and the verbose LDP Guides, lies the world of HOWTOs which do exactly as their name suggests. HOWTOs and their Mini counterparts (which are

shorter tutorials on more specific issues) contain detailed guidance on various aspects of configuring and using Linux. They range from, for instance, the Printing HOWTO, which covers printing issues in general to the Print2Win mini HOWTO which details how to set up your Linux box to output to Windows print share.

HOWTOs vary in both style and substance. The USB Digital Camera HOWTO, to take a random example, provides details not just of how to get your camera up and running, but covers what is going on behind the scenes when you run the included script, so you can check things are happening as they should, and offers a selection of troubleshooting tips if things go wrong.

HOWTOs are written in SGML, but are available as plain text, Postscript, PDF or HTML files. They are usually released under the GNU Free Documentation License which allows

redistribution and modification as long as the altered text is released under the same license. However, it is becoming increasingly common to write HOWTOs using the Wiki system of community collaboration where anyone can modify the text by simply hitting a link at the bottom of the page.

LXF Forum & LUGs

Sometimes it can seem there is nowhere to turn and every possible avenue has been exhausted, especially if your problem appears to be too specific for the general guides or HOWTOs. Firstly, you could write into *Linux Format's* Answers pages. As the magazine has a longish lead time and a lot of requests, your question may take a while to be answered, but the solutions tend to be thorough and well researched. The other alternative is to visit the magazine's website which has a number of sections dedicated to technical queries, hardware issues, programming languages and general help. As well as being able to call on lxfadmin (when there's no pressing deadline) and the ever present 'anonymous', the forums are also frequented by the likes of Jeremy, Nelz, Fingers99, Rhakios, Erin and many others with knowledge and experience of using Linux in various situations.

Of course, Linux user experience isn't limited to the LXF forums. User Groups often have mailing lists to supplement their RL activities (check out the LUG pages), www.lug.org.uk, www.linuxforums.org and www.yourlinux.co.uk, among many others, have extensive forums and

"If you've exhausted the paper and online manuals, the next stop should be the website of your distribution vendor."

entities such as IBM, Mozilla, Sun and OpenOffice have support systems in place to answer questions and help you sort out your problem.

We hope to have shown you by now that the support mechanisms for Linux are many and varied, and if you still haven't solved your difficulties, it might be time to consider hiring a consultant to help you out. But that's another topic all together... [LXF](http://www.lxf.co.uk)

Do your own documentation

Writing and printing

Documentation is one of the less glamorous aspects of open source development and, as such, is often left to the last moment (if it gets done at all). But it's the perfect job for users who'd like to contribute to a project without having to learn another language. The big projects are always looking for users (who, after all, experience the application on a day-to-day basis) who can string a sentence together to help compile online help, tutorials and HOWTOs.

The real benefit here is that once you've 'been through it', pulling together all possible resources into a coherent whole, you will have a unique perspective on an issue from which many others could benefit. It would be a shame to waste it...

If you find that you're having to print a lot of bulky docs, get some good quality paper, and print on both sides. The most straightforward solution for PDF files is to print odd-numbered pages first, then reload the printer with the odd pages, and print the even-numbered pages on the other side. To save tinkering with printer settings too much, Plain Text files can be most easily printed on both sides of sheets by pasting them into your choice of word processing app or text editor that allows page breaks. Remember to check the required orientation of the paper first by doing a single-sheet test to avoid your documentation coming out overprinted or upside down.

What on Earth is... SVG?

Due to the limitations of the Bitmap format, scalable images for Internet use used to require huge amounts of bandwidth. Macromedia Flash partially solved this, but it's still a proprietary format. **David Coulson** outlines the XML alternative.

» It seems a bit extravagant having a whole section of the magazine on a graphical file format. I don't recall seeing What on Earth is a JPEG in any past issues of *Linux Format*. Why?

Good point. But – and there has to be a but – Scalable Vector Graphics (SVG) is more than just another file format. For a start, it has been designed from the ground up for use on the Internet, which means a lot of emphasis has been placed on creating a format capable of sophisticated results but keeping the file sizes to a minimum.

It has also received the blessing of the World Wide Web Consortium (W3C), joining the likes of HTML, XML and ECMAScript as a 'standard' format. This ensures that no one can 'own' the format or restrict it, for instance, to a specific system architecture or application.

As the name suggests, SVG graphics are scalable due to the fact that they are 'vector' images made up of mathematical descriptions of objects rather than individual pixels. In the case of a picture featuring a plain red circle 140 pixels in diameter on a white 150x150 pixel square, a bitmap image file would need to house information on the Red, Green and Blue (RGB) values of 22,500 pixels. JPEG compression (which relies on measuring – and losing – adjacent pixel values) can reduce file sizes significantly, but really good compression ratios often lead to jagged edges and dirty colouration.

The same image saved as an SVG file would simply state that the viewer (SVG capable application) should draw a 150x150 pixel white square with a red circle 140 pixels in diameter in the centre. The browser or other SVG viewer takes care of actually rendering or rasterising the


image, automatically anti-aliasing between the white of the background and the red circle to create clean lines. Moreover, the same image could be configured to scale with the viewer window, making imagery degrade a lot more elegantly than could be expected with a bitmap, which tends to become indistinct as resolution drops and jagged as it rises.

The downside in the adoption of SVG is a slight performance hit which is introduced as the browser has an extra stage (rasterising) in the display process that is not needed with pre-rendered graphics. It wouldn't appear to be a major hurdle in this high-powered era though, as most of the average desktop's processing power is left largely untapped.

» So it's a way of displaying images on the Internet as vectors rather than bitmaps then. Just one word for you: *Flash*. What's the point in having yet another graphics format?

There are similarities. Both *Flash* and SVG are structured (vector) formats, designed for the World Wide Web and feature a high degree of interactivity;

```
<svg width="12cm" height="4cm">
  <desc>Example circle01 - circle expres
  <circle cx="6cm" cy="2cm" r="1cm"
    style="fill:red; stroke:blue; st
</svg>
```



we'll come to the file format characteristics in a moment, but there is a more fundamental difference between *Flash* and SVG. As it is estimated the some 95 per cent of web users (across all popular platforms) are able to read *Flash*, adding a competing graphics format would appear to be an endeavour of Betamax-esque foolishness. There are two things that mitigate against that assessment:

- *Flash* is a proprietary format owned by Macromedia. While the company currently develop their *Flash* browser plugin for *Mozilla*/Linux (which in turn allows *Konqueror* to see *Flash*), it is not really prudent to rely on the largess of one company to access an increasing number of websites, especially when that company is attracting the attention of Linux-hostile predators within the industry. Think *Apple Quicktime* or *Windows Media Player*!

- SVG is a W3C standard. At present you'll need to download a fairly substantial plugin (if one is available, see below) or use a specially configured *Mozilla* build to see anything (and even then your mileage will vary), but future browsers and operating systems will support the standard as a matter of course. Adoption of a technology appears to be directly related to its availability to end users, so SVG easily has the potential to become as ubiquitous as gif or Javascript.

When it comes to file differences, the biggest is that .swf is a binary format and .svg is text-based, in fact it is built on XML. While this means that file size for the latter tends to be slightly larger, the result is

human-readable and so can be viewed and altered in a text editor. *Flash* can deal with some XML data – via XML-Socket and XML-Object, but only SVG is capable of embedding foreign XML namespaces and self-defined attribute data. It is also CSS and XSL compatible, fully DOM compliant and can handle XSLT transformations.

As SVG has grown out of XML, it makes it the perfect format if you need to generate graphics – for instance maps, graphs or charts in particular – from live XML data.

» Everything's XML these days isn't it? How can a graphical format be text-based, it seems a bit implausible?

It would be if we were talking about pictures, a photo for instance, made up of individual pixels. But we're not. The pictures you would create with SVG will be based on primitive shapes, bezier curves, fills and line styles, just the sort of thing you'd draw with *Kontour*, *OpenOffice.org* or *CorelDraw*.

For instance, let's say that an idealised representation of a beach at sunset consists of a rectangular background with a graduated tint going from red to yellow; an orange circle on the right of the image for the sun; a dark blue rectangle for the sea; and a dark yellow rectangle for the beach. You could even add a 50 per cent transparent, white four-sided shape running from horizon to beach to suggest the sun reflecting off the water.

What we have is a human user-readable

description of a scene, and creating an SVG version wouldn't involve an enormous amount of work (or generate reams of code).

There are big benefits to being text-based as well: pages can be indexed and searched by search engines, screen readers (used by blind surfers to access web pages) can be easily accommodated and users get access to the source, as with HTML.

» What does a typical bit of SVG code look like?

As with any XML document, an SVG file begins with a version declaration, which is followed by a document type definition:

```
<?xml version="1.0" encoding="iso-8859-1"?>
<!DOCTYPE svg PUBLIC "-//W3C//DTD SVG
20000303 Stylable//EN"
"http://www.w3.org/TR/2000/03/WD-SVG-
20000303/DTD/svg-20000303-stylable.dtd">
```

The next section, in this example at least, deals with the outputting the image to screen, first defining its position and then setting the quality and precision of the rendering.

```
<svg xml:space="preserve" x="-3.65914in" y="-1.01866in" width="3.31229in" height="3.31229in"
style="shape-rendering:geometricPrecision; text-rendering:geometricPrecision; image-rendering:optimizeQuality"
viewBox="-3659 0 3312 3312">
```

One of the real benefits of SVG is that you can embed Cascading Style Sheets within the document »



WhatOnEarthSVG

« (or refer to a remote CSS file). We've defined three styles here, one stroke and two fills:

```
<defs>
<style type="text/css">
<![CDATA[
.stro {stroke:#000000;stroke-width:3}
.filo {fill:none}
.fil1 {fill:#FF0000}
]]>
</style>
</defs>
```

Finally we have the actual description of the two elements which make up our image. In this case a rectangle which uses two of the CSS styles configured above (.stro and .fil0) and an ellipse that uses .fil1. The rectangle has the same width and height, thus making a square:

```
<g id="Layer 1">
<rect id="32270828" class="fil0 stro" x="-3658" y="2" width="3309" height="3309"/>
<ellipse id="32270300" class="fil1" cx="-1978" cy="1706" rx="1306" ry="1306"/>
</g>
</svg>
```

Note that the whole thing is enclosed in <svg></svg> tags. While HTML is often forgiving of the occasional missed closing tag, XML is a stickler for the rules, and will just throw up an error. Where we don't have full tag sets – in the image definitions, for instance – we've used a unified open/close tag <ellipse id... />, which is valid.

On the subject of validation, once you've created your images, why not visit <http://validator.w3.org/>, input the URL of your site, or local file address and make sure it passes muster? It's something we at

LXF couldn't do without; the site has basic and more advanced validation tests and will also make certify that your HTML, XHTML and any other standard mark up language you've used is correctly coded.

» ...and what about the final images?

Viewed on a suitable application, the above code will render a red circle in the centre of a square drawn with a three pixel wide border. Simply adding **stro**, the **class=""** element in the ellipse definition would give the red circle the same border as the square, or we could change the colour of either element by adjusting the hex numbers in the style definition.

» To the untrained eye, it would seem something of a backward step to have to laboriously write out my image in Emacs?

We won't argue there. Fortunately you won't have to take such retrograde action. The image associated with these examples was created in *CorelDraw* and simply exported as an .svg file. It can then be opened in any text editor for examination or fine-tuning. SVG import and export are now regarded as vital for any graphics package, and in the Linux world *Kontour*, *Karbon 14*, *Sketch*, *OpenDraw* and *Sodipodi* have support for the format. In fact, there is a strong argument for adopting it as the native file format stressing, as it does, the open, collaborative ethos that these applications are built on.

However, each of these applications seems to output their SVG in quite different ways and some are more efficient and comprehensive than others. For instance, the image mentioned above, when exported from *OpenOffice.org* defines the ellipse as a

polygon with over 200 points, rather than relying on the more prosaic 'ellipse' definition.

Still, it is early days yet, and all these applications can open and save the files complete with gradients, transparency, styles and text.

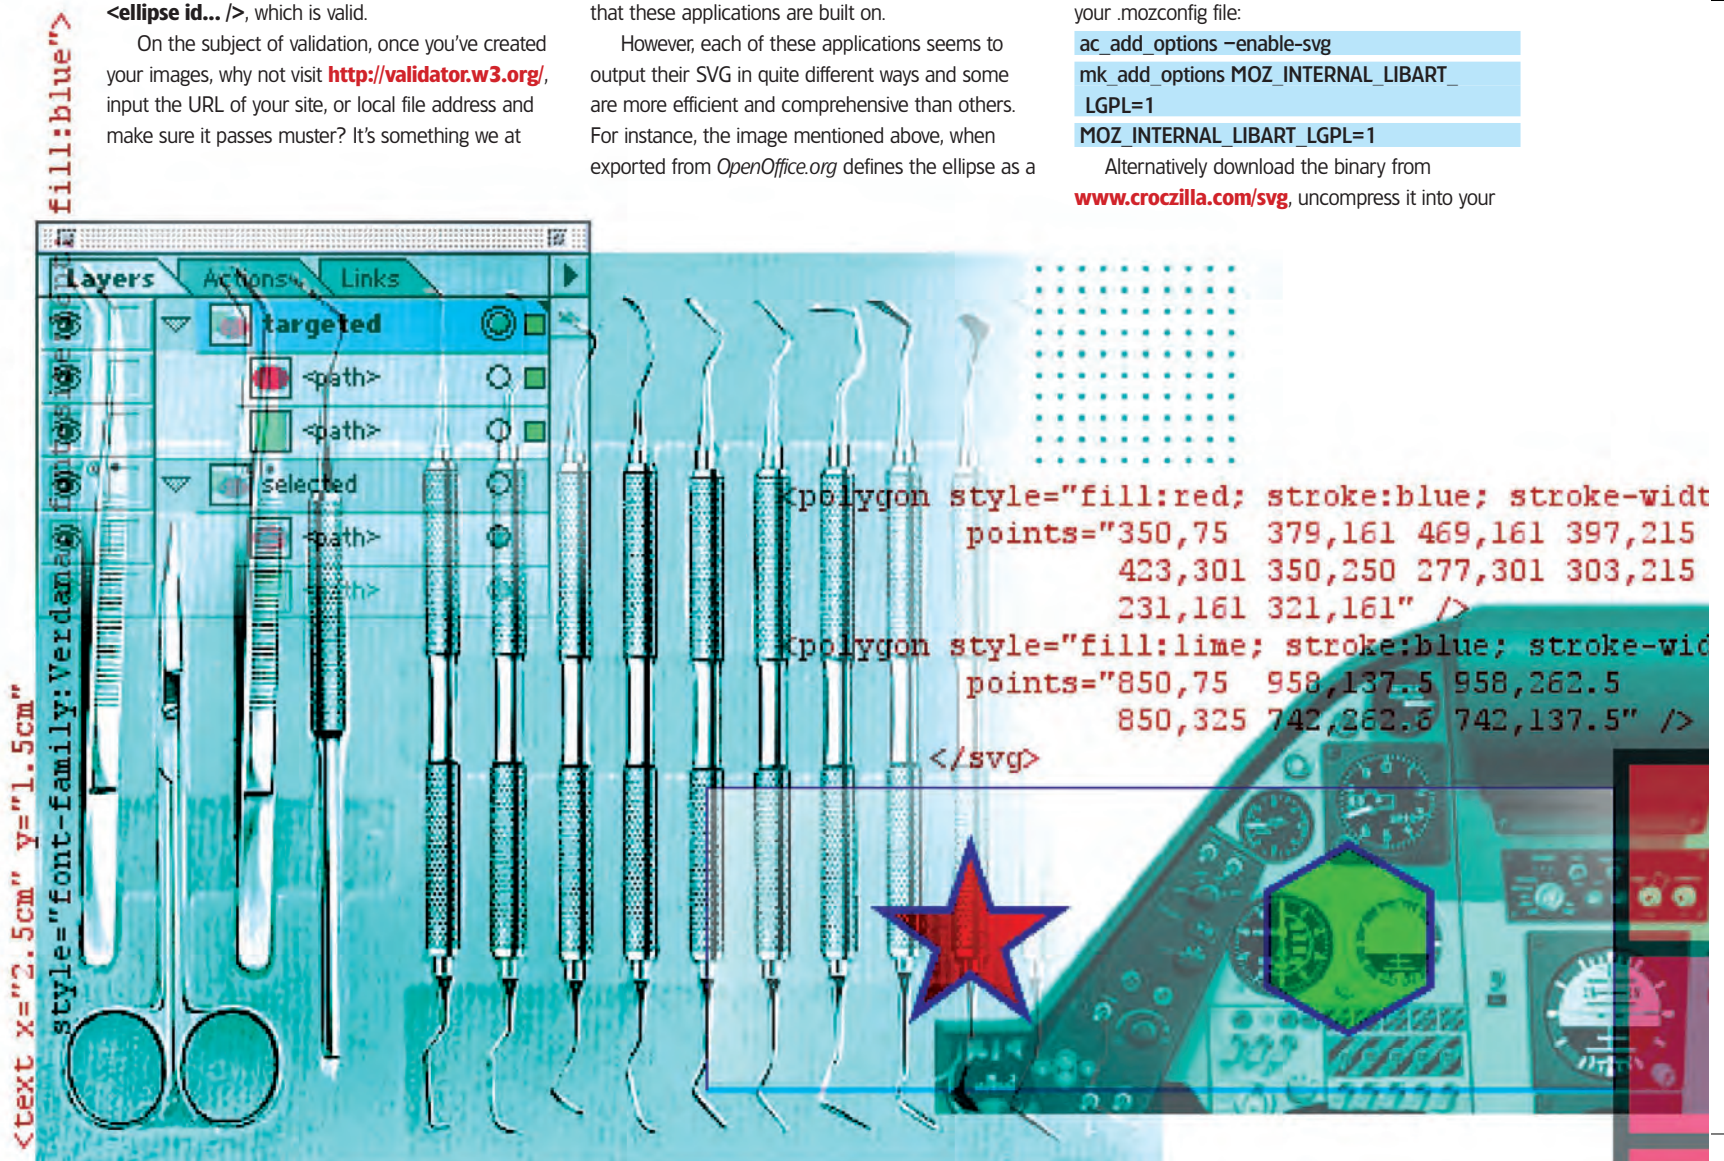
» That's a relief! But once I've made my pretty pictures or whatever can I look at them from any browser?

This is where – as of January 2003 – we fall over slightly. The most common viewer was created by Adobe (who are partly responsible for the SVG standard) and used to work with *Mozilla*. But as the browser project edged towards its 1.0 release, compatibility with the plugin was broken and – as Adobe had used a 'non-frozen' API to create their product – it has yet to be updated. Mozilla.org regard the issue as an 'evangelism' problem and encourage users with problems to contact Adobe to fix their side of the applet.

This is not an insurmountable problem for *Mozilla* as a plug-in solution for SVG was always regarded as a stop-gap measure. And, indeed, you can either compile a *Mozilla* build with SVG support or download a pre-built binary. The implementation isn't comprehensive yet and seem a little temperamental on some of the sites we've tested, but it does work and it will be included in the main browser in the near future. You can pull and build *Mozilla* with SVG support by adding the following to your .mozconfig file:

```
ac_add_options --enable-svg
mk_add_options MOZ_INTERNAL_LIBART_
  LGPL=1
MOZ_INTERNAL_LIBART_LGPL=1
```

Alternatively download the binary from www.croczilla.com/svg, uncompress it into your



home directory and double-click on the *Mozilla* icon. The version on croczilla is currently 1.3a, but Alexander Fritze seems to be pretty close to official releases with his work.

KDE/Konqueror has been on the cusp of gaining real SVG support for a while (at the moment it can use KSVG2PNG to render files on the fly), but even the latest 3.1 release fails to come up with the goods. Oddly, static SVG support is available in Qt but, according to one KDE developer, the results are 'ugly as hell'. This state of affairs led to the creation of the KSVG project which intends to add support to the *kparts* armoury, thus making it available to all native KDE applications. It was due to go into the main distribution in 3.1, but looking at LXF's copy, it doesn't seem to have made it.

GNOME is a slightly different proposition in that, via *Nautilus* it already has some limited SVG support (see below) but the *librsvg* SVG rendering library and SVG-gtk engine are now being co-maintained by Dominic Lachowicz, the author of AbiWord's SVG plugin, and he plans to keep GNOME at the forefront of W3C standard support. GNOME looks like it may lead the way here. One bizarre thing: *Nautilus* can render SVG for its own icons, but it can't manage it when displaying web pages thanks to its reliance on *Mozilla* for that task!

» Are Adobe likely to reinstate Mozilla support with their SVG viewer application?

Adobe don't seem overly concerned to update their plug in for post-1.0 *Mozilla* releases – though their development arm is famously unwilling to talk about upcoming releases – suggesting that native support will negate the need for a plugin altogether. Still, depending on your half empty/half full assessment of the situation, it might be a good idea to email and ask them to sort it out. Demand is the only way to change their minds.

» So what are the alternatives?

Well, there are alternatives. *Batik* (<http://xml.apache.org/batik/>) is a Java-based

toolkit which forms the basis of projects such as Luxor XUL and XML_SVG2image, a PHP class which creates JPG images from .svg files.

For browsing and/or authoring the W3C's own project Amaya offers an incomplete implementation of the standard, though this should be updated in the very near future.

» It sounds like the ideal format for building desktop icons: resizable, clean, small footprint etc. Will SVG be implemented on the KDE/GNOME desktops of the future?

As mentioned above, KDE is on the cusp of having its default Crystal icons available as SVG files, but currently its stuck using pre-rendered PNGs. GNOME, on the other hand, can do it, and do it well. The only shortcoming is the number of themes designed to make the most of it. There's one and it is called Scalable Gorilla (SG). Developed as part of the Ximian desktop, SG is a pretty cool icon set capable of being used at the highest resolutions, while still looking good on a bog standard desktop. The latest version shipped without standard *Nautilus* icons (navigation etc) which was a bit weird, but if you're using GNOME above 1024x768, they make a great deal of sense.

» The Internet is primarily an interactive medium. Is there anyway SVG can be used to make my images more dynamic?

SVG can do pretty much all that *Flash* can do (and more besides), so the simple answer is yes! There are so many possibilities but, briefly, you can have mouseover/click/off events or timed events; transformations including rotate, flip, scale, skew; clipping paths; variable transparency – including a group transparency tool; tons of different line styles, line caps, dashes; even more fill styles (including vector patterns); embedded ICC colour profiles; *Photoshop*-style filters and frame, time, interpolation or script-based animation. Oh, and these animations can feature changes in colour, position, opacity, geometry and position/orientation along a path. In fact, almost any element can form part of an animation, if you want it to.

Finally, SVG should be your first choice is you're generating on-the-fly imagery or animation from GIS applications such as SiCAD, MapInfo or a database source.

So, in a nutshell, you can use SVG in tandem with other web standards such as Javascript and DOM, to create graphical user interfaces, web applications, personalised pages, graphs from back-end data, and real text. And all this in a standard format which, in the future, will be supported by all the major browsers. Oh, and we mentioned that it was completely Open didn't we? Thought so.


» Blimey, I'm impressed! Are there any decent examples of what SVG is capable of on the Internet?

There are tons of sites dedicated to showing off the capabilities of SVG, just be prepared to see the 'SVG tiger' an awful lot. Initially, if you're checking out *Mozilla*'s SVG work, the best place to start is www.croczilla.com. Here you'll find some useful demonstrations of the various ways you can manipulate graphics using the format. For a little more interactivity, Adobe have a series of demos on their site (www.adobe.com/svg/) the most impressive of which is a small graphics application built using SVG, that is itself capable of outputting valid SVG files. A nice touch, but it will put a bit of a strain on your browser, even if you're using Adobe's own plugin.

One of the areas where SVG will have a major impact is in online cartography and www.carto.net have some good examples of the way it may be used in the future.

If you're interested in the future of the Linux desktop, check out the open source SVG icon project (<http://svgicons.sourceforge.net/>) which aims to produce a standard set of icons for use on both KDE and GNOME. The set currently includes some 120 icons. The other icon development to keep an eye out for at present is Everaldo's CrystalSVG set which is 'about 80-90 per cent complete' and should become the standard set for KDE in the very near future.

» Where can I go for more information?

The first place to look is the W3C website. Here you'll find the official standard definition as well as documentation and an examination of all the available elements: www.w3.org/TR/SVG/. A complete reference to all standard SVG elements (in a rather more readable format than the W3C site), with an amazing number of examples: www.zvon.org/xxl/svgReference/Output/#. The SVG Foundation have a number of features available discussing potential uses of SVG in areas such as interface design and web services, which may inspire you to think of applying SVG in new ways: www.svgfoundation.org/. The official Mozilla work resides at: www.mozilla.org/projects/svg/. KDE's SVG-related info is at: <http://svg.kde.org/>. GNOME users can find information at: www.levien.com/svg/ 

width:10"
215
215
-width:10"
</>
maps
logos
interfaces
animations
illustrations



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 suggestions in our forums? www.linuxformat.co.uk. Hope to hear from you soon!

Nick Veitch EDITOR

THIS MONTH TEACH YOURSELF...

Mozilla >>

PART 2: Manipulating the sidebars and the throbber for that extra-special personalisation **p68**

The GIMP

If you need an image editor as capable as *Adobe Photoshop* but for Linux – start here! **p72**

PHP


Document Distiller isn't the only way to create PDF files – you can use PHP too! **p76**

Python >>

"Short, concise, readable classes that express a lot of action in a small amount of clear code" Guido Van Rossum **p80**

Disk backups

Save entire HD images and protect your data with Partition Image – step-by-step backup tutorial. **p84**



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.

TIP OF THE MONTH!

Winston Churchill declared "History will be kind to me, for I intend to write it".

You are making history every time you type a command in the shell, and with this quick tip, we prove it can be useful. In many common Linux shells, the history command is enabled by default, so you won't have to do anything to start using it. Open up a shell and type:

```
history
```

You'll get a numbered list of all the commands you have entered which stretches back as far as the history buffer will allow (usually about 500 entries). This serves some useful purposes in itself. You can see how

A question of History

exactly you constructed a tricky command, or if anyone has been logging on as you. But history is more than just looking back at what's gone before, it's about repeating things too.

You've probably used the console history function before. A cursor-up press will replace your current input line with the last command executed. That's all very useful, but there are other tricks which can speed up your use of the contents of the history buffer. One of the most useful is the **!!** construction. This simply tells the shell to use the last element of the previous line. Suppose you edited a text file and

then wanted to do a quick word count before copying the file into some distant directory as a backup:

```
vi thisismytext.txt
```

```
wc !$
```

```
cp !$ ~/textfiles/!$.bak
```

Execute these commands, and you'll see that the command line is repeated with the substitutions, then any output is listed. If you use the cursor key-up trick to look at the commands you just typed, you'll see that the history buffer includes the full, substituted command.

There's more tricks to manipulating history. Check out the documentation of your shell for more info.



TWEAK YOUR BROWSER

Mozilla: throbber and sidebar manipulation

PART 2 After skinning the UI last issue, **Andy Channelle** demonstrates how to provide a little user feedback and brand awareness to your distro with the throbber and also add functionality to a website using *Moz's* unique Sidebar.

So far we've stressed *Mozilla's* credentials as a tool for branding; adjusting the user interface to, perhaps, mirror a corporate colour scheme or style and integrate the suite into the look of an intranet. This month we'll continue that theme and dissect the *Mozilla's* throbber to create a short animation with *The GIMP* which will be activated when the browser or mail client is busy. As well as providing vital user feedback, the Throbber can be another useful branding tool, although getting a compelling message across in a 32x32 pixel square over 30 or so frames is not a walk in the park.

Sidebar

The sidebar is the space on the left edge of the browser which acts like *Internet Explorer's* explorer bar on steroids, it is basically a browser window/frame capable of displaying any web content and interacting with the main content space. As with the rest of the browser, the sidebar can be skinned using a combination of CSS and graphics.

Sidebar pages are organised using 'tabs' that can be moved, deleted or configured in a number of ways. The default selection – depending on which version of *Mozilla/Netscape* you're using – may contain a set of developer-centric options or a basic

search/favourites/what's related selection. To change this, simply click Tabs and select Customize Sidebar. This will open a dialog containing the currently available sidebars on your computer. If nothing here inspires, select Sidebar Directory to visit Netscape's selection of additional sidebars, containing everything from technology news and programmers' references to sticky note applications and e-commerce sites.

In keeping with the possibility to "change everything", we're not stuck with the sidebars that Netscape, CNN or anyone else hope to 'push' to our desktop, building our own, and then making it available to the rest of the great unwashed is a pretty easy job. Creating a Sidebar is essentially the same as designing any web page – though there are a few simple distinctions – and can be done with any standard design tool such as *Mozilla's* Composer, *Quanta* or *Bluefish*.

By default the standard *Mozilla* sidebar is 162 pixels wide, though if we factor in the scrollbars, any webpage or application will be limited to a paltry 144 pixels. While this is a big restriction, it does have a number of benefits, not least the need for brevity and simplicity in our design which, in this case, is a forum browser for the *Linux Format* website. Using Composer, we've pulled together various elements including a small graphic, a title and a selection of links copied directly from the *LXF* site. In raw code the link for the 'help' section looks like this:

```
<a href="http://www.linuxformat.co.uk/modules.php?op=modload&name=phpBB_14&file=index&action=viewforum&forum=1" target="_content">Help!</a>
```

The first section follows the usual HTML format for links, but the important bit to note is the target which is set to `_content`. This ensures that whenever we click on a link in the sidebar, it opens the resulting web page within the main window. Leaving the `target=` tag blank will mean the new page opens within the sidebar. It is exactly the same method used in multi-frame websites. As with normal web pages, anything can be a link, so the *Linux Format* logo links to the front page of the LXF website, the forum links go to the appropriate location, while the "Subscribe Here" link will take the viewer directly to Future Publishing Ltd's secure subscription site.

As an aside, banner advertisements, marquee text or other flashy animations are definitely a bad thing when it comes to sidebar design, and putting a popup or 'blinking' text, though not officially regarded as a criminal offence, will ensure your handiwork remains largely unloved and unused. If you must advertise – and the sidebar would appear to be a good place to do so – make sure it is unobtrusive, preferably plain text and linked to a more elaborate sales site.

Our initial design called for a drop-down jump menu using a smattering of Javascript (similar to that seen on the main LXF forum page) but a scan through the original sidebar development documents suggested this method is frowned upon for the performance hit Javascript may impose on users. Size here is important and our browser comes with an initial download of under 10KB, though once the image is cached, that will come down to just over 4KB for updates.

When it comes to specifying the period between updates, we can add a simple `<META>` entry to force the page to reload at specific intervals, in this case once an hour:

```
META HTTP-EQUIV="Refresh" CONTENT="3600;
```

```
URL=http://[your_url].html"
```

This instruction will be called for and the update will take place as long as the tab is left open.

Add a sidebar to your site

There's no point in spending our time creating a sidebar tab unless it gets released into the wild, so once we've finished coding, we need to put a link on a web page to add the tab to a viewer's browser. We do this by calling the `addPanel()` function; simply adding the following short Javascript to the body section of our web page, with the tab name and URL inserted within the brackets in line four:

```
<script>
function addsidebar() {
if (window.sidebar&&window.sidebar.addPanel)
    window.sidebar.addPanel("LXF Sidebar","http://andychannelle.
    pwp.blueyonder.co.uk/lxftab.html","");
else
    alert("You need Mozilla (or Netscape) to use this particular
    feature. Visit www.mozilla.org")
}
//—>
</script>
```

The third option in the `addPanel()` function (left blank here, though it's vital to keep it in) would be used if we wanted to give users the ability to customise our sidebar.

Once the actual function has been defined, you can add a link or graphical button to launch the sidebar in the same way that you would link to another page, substituting the URL for the Javascript function:

```
<a href="javascript:addsidebar()"></a>
```



Small changes

Tailoring Mozilla profiles

There are a number of ways you can change *Mozilla*'s look and feel. We've been concentrating on the creation of a browser suite, but if there are just a few elements of the application that don't look or function as you'd like, you can hack into a couple of easily available files and change many aspects of it.

The great thing about this way of working is that your changes are limited to the profile/chrome directory in which the configuration files are stored. So your single *Mozilla* installation can be tailored specifically to meet the needs of your whole family/workgroup with just a little CSS and Javascript.

The first thing we'll need to do (if it has not already been done) is create a named profile. Open the console and `cd` to the *Mozilla* program directory:

```
cd usr/lib/mozilla
```

Now type `mozilla -profilemanager` to fire up the application. Select 'create new profile' and follow the directions given by the wizard.

By default, *Mozilla* will store your profile in `~/mozilla/[Profile Name]/[random string].slt/`. Navigate here and take a look inside the `/chrome` folder, where you'll find examples of both `userChrome` and `userContent` to follow.

By the way, if you browse the *Mozilla* help for information on the Profile Manager, you should note that it – like much of the official help – is quite obviously Netscape-specific and may not bear any relation to the actual directory structure of your installation.

USERCHROME

This file changes the look of the browser. You could, for instance shift *Mozilla*'s sidebar from the left to the right-hand edge of the screen:

```
window > hbox {
    direction:rtl;
}
```

```
window > hbox > * {
    direction:ltr;
```

or change the font size and family used in dialog boxes:

```
window {
    font-size: 3.5mm !important;
    font-family: helvetica !important;
```

USERCONTENT

`userContent.css` can alter the way *Mozilla* displays web content on any site. For instance, the blinktag, introduced in Netscape 1.1, is an aberration. Fortunately you can make sure your screen is never corrupted by its eye-scorching foulness by adding the line.

```
blink { text-decoration: none !important; }
```

Less annoying are instances of marquee text, but you can still halt this in its tracks:

```
marquee {
    -moz-binding: none; display: block; height: auto
!important;
```

You could just use `display: none !important;` method for marquee text, as we've done with blink, but this way you still get to see the contents of the

marquee. You should also remember to add the `!important` element to over-ride the default settings.

USER.JS

Preferences selected in *Mozilla* preferences dialog are stored in a file called `prefs.js`, and this is overwritten whenever you make changes from within the browser. However, there's more options available than you'll find in **Edit>Preferences**, the problem is that adding them to `prefs.js` will only last until the next time that file is written to. The solution is to change settings, add comments etc into a new file in the same directory call `user.js`. This file is read by *Mozilla* but never written to.

In this file you can use something like:

```
user_pref("capability.policy.popupsites.sites",
    "http://www.adverts.com http://www.popunders.com
    http://www.someotherurl.co.uk");
user_pref("capability.policy.popupsites.Window.open",no
    Access");
```

to selectively kill popups if, for instance, you regularly surf sites that actually have useful popups rather than irritating advertising that you still want to be able to see and click on if necessary. Finally, if you want to make doubly sure `<blink>` doesn't work, you can also switch it off here:

```
user_pref("browser.blink_allowed", false);
```

Further examples of the options configurable in these files can be found on the *Mozilla* website.

www.mozilla.org/unix/customizing.html

Moztech 101: part two

A brief look at some of the underlying technologies that make *Mozilla* so flexible

While CSS (see part one, *Linux Format* issue 37) defines the way *Mozilla* looks, Javascript is mostly concerned with telling it how to act.

Javascript is a scripting language developed by Netscape to add interactive elements such as roll-overs, drop-down menus and time-sensitive widgets (clocks etc) to bog-standard HTML pages. The actual program is contained within a `<script>` tag and can be called from a central, either remote or local, location.

Basically, Javascript deals with objects which commonly have properties and methods. One of the most common objects you'll encounter is 'document' which controls the way pages are displayed. So, in:

```
<script language="javascript">
<!--
document.write("This text will be displayed on the
page")
```

```
//-->
```

```
</script>
```

document is the object and **write** is the method. We could augment this short program with a number of properties such as **bgcolor**, **fgcolor** and **lastModified** which alter the background and foreground colour respectively and print out the date and time the page was last modified. For instance:

```
<script>
var dateandtime
dateandtime=document.lastModified
document.write("This page was last modified: "
+dateandtime)
</script>
```

The above script would output:

```
This page was last modified: 01/23/03 11:09:22
```

Mozilla has two Javascript implementations; the first -

Spidermonkey - is written in C and is based on the initial work down by Brendan Eich at Netscape, while the second - *Rhino* - was created in Java by Norris Boyd. Both are compliant with ECMA-262 Edition 3 and should rightly be called ECMAScript.

While most Internet users will be aware that Javascript is used primarily for displaying web pages, Javascript is also used in the configuration of the browser, just look for files ending with a .js extension and have a poke around (after backing up of course). A prime example is the prefs.js file which is updated whenever you use the preferences dialog in *Mozilla* (see earlier in this article).

You can find more Javascript information than any individual could possibly remember at http://www.objenv.com/cetus/oo_javascript.html which claims over 18,000 relevant links.

On clicking the image, our viewer is presented with a dialog asking if they want to install the sidebar. If you wanted to do a little *Mozilla* advocacy, you could substitute the **alert()** function with a dialog giving users of other browsers the option of upgrading and linking them with mozilla.org.

```
{
var rv = window.confirm("Sidebars are a Mozilla/Netscape feature."
+ "Would you like to upgrade now?");
if (rv)
document.location.href = "http://www.mozilla.org";
```

Finally, if you've yet to set up your own web page but still want to preview a sidebar, you can upload it to your host/ISP and then

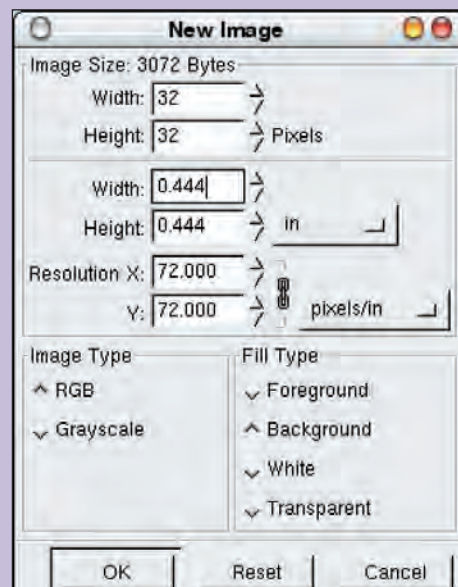
visit <http://wp.netscape.com/sidebar/preview.html>, where you can enter the tab name and URL and have the tab added to your own sidebar.

Adding a sidebar to your suite


If you intend to distribute a browser or theme complete with sidebar tabs included, the method for adding it is obviously going to be a little different, but only slightly more complicated. So far we've only viewed and altered .css and graphical files, but we must now venture into our first Resource Description Framework (RDF) file which, according to the W3C, "provides a lightweight ontology system to support the exchange of knowledge" on the Web.

The *Mozilla* throbber


Make your own adjustments to those 1024 pixels



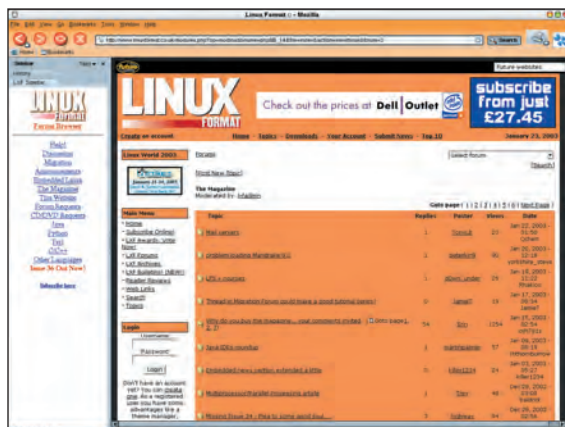
1 Open *GIMP* and create a new 32x32 pixel image. This will give you the first frame of your animation which can also double up as the 'resting' icon (the icon that is shown when the browser is not busy).



2 Now create your miniature masterpiece. You may discover that the art of icon design is not as simple as it seems. As your space is so small, complexity and detail is going to get lost, so think bold and simple, with high contrast colours and aim for an integrated feel with the rest of the browser. You also need to bear in mind that your image will be the first frame of an animation, so think which parts you intend to move. If you have an element that will be constant - say some text on a changing background - make sure it's drawn on a new layer. We've misappropriated the Lycoris logo (www.lycoris.com) here because it's simple, elegant and has potential as an animation.



3 Once your initial image is perfect, you can begin building the animation. The standard *Mozilla* throbber with the famous little red dinosaur is 30 frames in length, but obviously you can opt for an many or as few frames as you like or your animation requires. We've duplicated the first layer, containing the Lycoris flower, and then fiddled with each one to - hopefully - create a smooth animation. Each layer will become a frame when we export the image.




To add or remove panels, we need to navigate to /defaults/profile (or the named profile directory if you're setting up users) and open the 'panels.rdf' file in a text editor. The file is basically a list of panels available in the standard distribution, and we need to add our own entry.

After the `<RDF:Seq>` tag we will need to add a resource description for the panel:

```
<RDF:li resource="urn:sidebar:panel:lxftab"/>
```

The next section of code concerns what is needed to configure the panels, so we need to add this new section:



```
<RDF:Description about="urn:sidebar:panel:lxftab">
  <NC:title>&sidebar.lxf-tab.label;</NC:title>
```

```
<NC:content>http://andychannelle.pwp.blueyonder.co.uk/lxftab.html</NC:content>
```

```
</RDF:Description>
```

Remember that as this is based on XML, every tag needs a corresponding closing tag. The first line is the resource name and should correspond with the **RDF:li** entry above. The second line is the title that will appear in the sidebar tab and the third line provides the link (which can be local or remote) to the sidebar. Save the file, restart *Mozilla* and you should have a new tab in there.

The sidebar is a great navigation element on, for instance, a large intranet project, taking the place of a standard navigation pane or Javascript menu setup. However, it really comes into its own when combined with something like RSS which can deliver a scrolling news feed to a users' sidebar. We'll try to cover that in the next tutorial.

The Throbber


The throbber has been an essential part of the Internet experience since the very earliest graphical browsers. Connections back then were so slow that downloading even a marginally complex web page could take so long it would convince even the most patient user that their PC had fallen over. So, ticking away up in the top right corner is a simple 32x32 pixel animation, handily reassuring viewers that the computer is merely busy.

Like everything else in *Mozilla*, the throbber is based on a standard web technology, in this case the animated GIF, and can be easily recast using *GIMP*, or any other graphics app capable of saving the format directly or compiling a series of images into a single file. See below for a step-by-step introduction. **LXF**

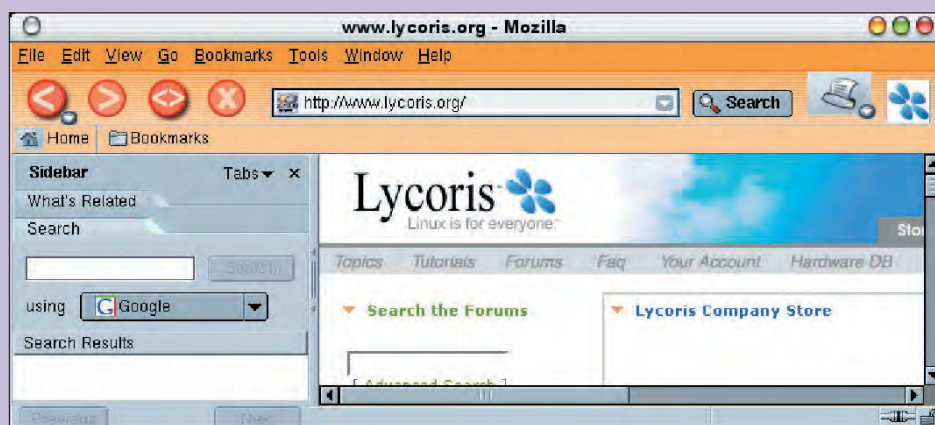
NEXT MONTH

XML

Next time we'll concentrate on adding a new toolbar to the browser and attempt to wrestle some dynamic content and a live news feed into the sidebar. We'll also begin the process of packaging up our browser modifications and looking into the XML technology which underpins so much of the *Mozilla* project and, it increasingly seems, the rest of the digitally connected world.



4 If you look into the /communicator/brand/ directory you'll find four gif files; there are two for the normal throbber and another pair for the small edition. Copy the first frame and paste it into a new file. Save this new image into /communicator/brand/ as throbber-single.gif. To make the 'resting' icon. Finally save your entire sequence into the same folder as throbber-anim.gif. Make sure save as animation is highlighted in the export dialog, and then set your animation speed (the default is usually fine). If you're skinning the whole application, you can then reduce your work to 16x16 and overwrite the other two files to create the 'full-screen' throbber.



5 Restart *Mozilla* and admire your handiwork.

The throbber is not the easiest thing to do well, but it is ever-present, and so is worth taking the time to get right to add that extra touch of professionalism to your version of the *Mozilla* browser. Check out some of the efforts of other theme builders at www.mozdev.org.

TutorialGIMP



ART AND IMAGE EDITING

The GIMP User Environment

PART 1 *The GIMP* is the grandfather of the Linux/GNOME desktop environment. **Michael J Hammel's** quick tour will get you started using this popular image editor.

The GNU Image Manipulation Program, better known as *The GIMP*, is the premier interactive image processing application for the Linux operating system. Its user interface is based loosely on the popular *Adobe Photoshop* image editing package. But while *Photoshop* provides a long, dual-column Toolbox with multiple pop-up menus, a cluttered menu bar, and an enclosed workspace, *The GIMP* provides a simplified Toolbox and menu bar, a free floating workspace, and sophisticated Canvas windows with tear-off menus.

The GIMP provides many of the same features as *Photoshop*, from layers and paths to blend modes and text manipulation. It also provides features not found in *Photoshop*, including support

for multiple scripting languages which makes *GIMP* perfect for web based animation production.

GIMP's flexibility and easy to understand interface, coupled with its sophisticated plug-in API, scripting languages and open source license has made it an integral part of the motion picture special effects industry.

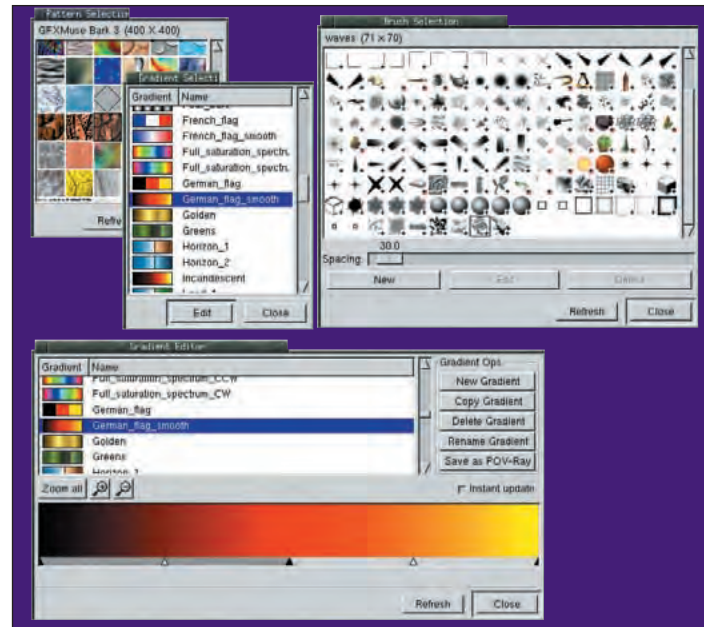
The GIMP interface

The main window for the *GIMP* is called the Toolbox. The tools should be familiar to *Photoshop* users, though the layout of the window is slightly different. First, individual tool buttons are not pop-up menus on their own. Double-clicking a tool button in the *GIMP* opens a Tool Options dialog which allows you to change the behaviour of the tool. This helps reduce the complexity of the tools by allowing them greater flexibility for a single tool instead of providing multiple tools that do similar things.

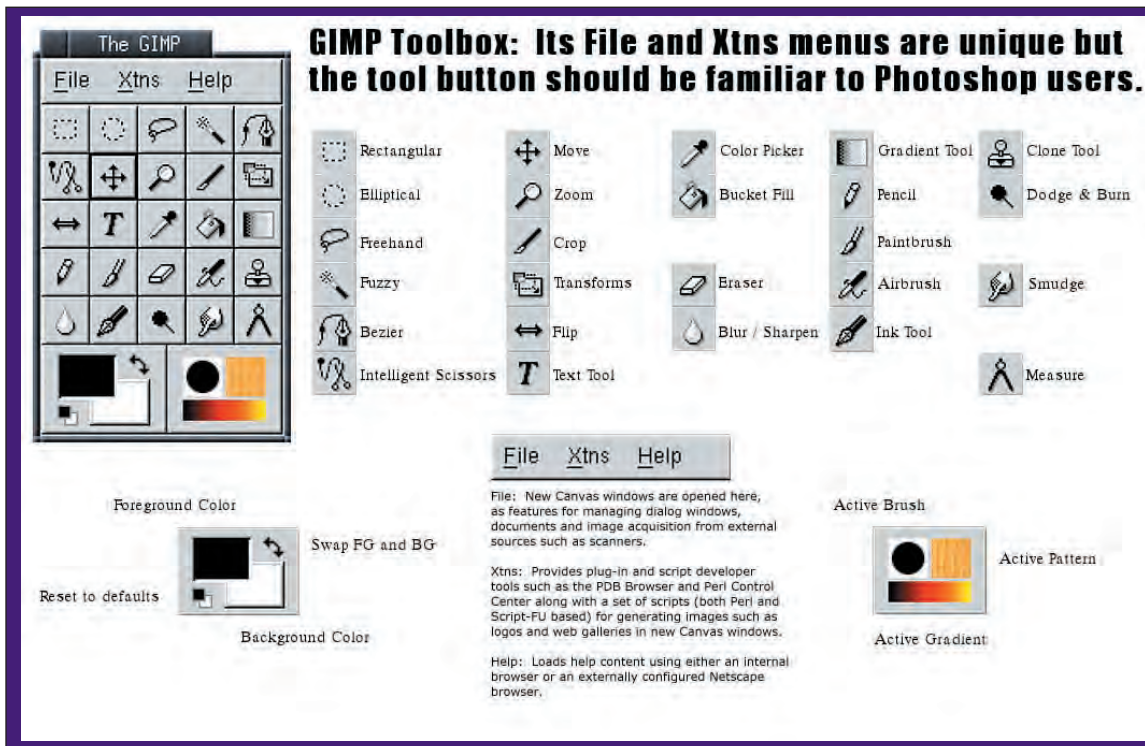
Next, the concept of Full Screen is irrelevant with the *GIMP* –



The working environment of *The GIMP*. Starting at the upper left: the Toolbox, a Canvas window, the Layers and Channels dialog, the brushes dialog, the Tool Options dialog and a tear off Filters menu. A typical *GIMP* session often has multiple Canvas windows opened, along with the Layers dialog. The brushes, patterns and gradient dialogs can be left open, or closed and quickly recalled using respective active buttons in the Toolbox.



The brushes dialog offers access to standard grayscale brush shapes, coloured brushes and Brush Pipes. Brush pipes offer artists brushes with multiple shapes, any of which can be imprinted on the Canvas window depending on direction, speed and even pressure applied by the paint tool. Predefined patterns can be easily saved using the '.pat' file format. Gradients can be created and edited using the Gradient Editor.



GIMP Toolbox: Its File and Xtns menus are unique but the tool button should be familiar to Photoshop users.

URLs

Online *GIMP* resources

GIMP:
<http://www.gimp.org>
GIMP Registry:
<http://registry.gimp.org/index.jsp>
Graphics Muse Tools CD:
<http://www.graphics-muse.com/gfxmuse/gfxmuse.html>
Tutorials and other useful links:
<http://www.graphics-muse.com/cgi/gmcat.pl?id=11>

since the workspace is open (not constrained by the application) it is possible to move and resize windows at will. Iconifying the *GIMP* will only affect the Toolbox itself, not the other windows. However, a quick way to hide all the windows is to press the Tab inside a Canvas window. This will cause *GIMP* windows to be hidden. Because of the unconstrained nature of the *GIMP*'s interface, there are no icons on show in the Toolbox for setting the application display size like the ones that inhabit the bottom of the Photoshop window.

Finally, menu options in the *GIMP* that create new images or modify existing ones can apply either to the active Canvas window or to a new Canvas window. The concept of an active Canvas window is sometimes hard to grasp for new users of *GIMP*. One of the most common questions is "Where are the Filters?" Each Canvas window has its own menu, accessible from the menu button in the upper left corner of the window or by holding down the right mouse button inside the window.



TutorialGIMP



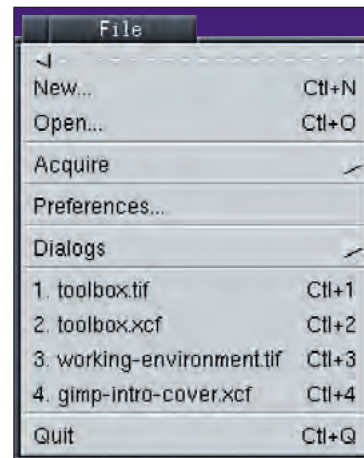
This is the catch-all menu for plug-ins that don't quite fit anywhere else. Many of these tools are script based, either Perl or Script-Fu. The most useful of these are not included in the stock *GIMP* distribution, such as the Web gallery tools and the animation and video tools. But most can be downloaded from the *GIMP* Registry.

The File and Xtns menus

GIMP's Toolbox provides three menu choices – File, Xtns and Help. While the purpose of Help is fairly obvious – links that launch web browsers to online documentation and other useful sites – the purpose of Xtns may not be. In fact, it's commonly asked how these menus relate to the Canvas menu.

If you're used to *Photoshop*, one of the first hurdles to migrating to the *GIMP* will be understanding the active Canvas concept. Menu options for any menu are either associated with the current active Canvas window or are independent of any Canvas window. This means that menu options that render new images will either do so in the currently active Canvas window or in a new Canvas window.

The Xtns menu provides options that are not dependent on any open Canvas window. The menu is divided into two parts. The top portion provides options that are more plug-in and script developer oriented. These include the DB Browser and PDB Explorer for searching for plug-in API information and the Parasite



The File menu is where you'll open existing documents or create new Canvas windows. You can also configure *GIMP* directories and various behaviours using the Preferences dialog. If you have scanner plug-ins such as XSANE or QuiteInsane they would appear under the Acquire submenu here.

editor for adding annotations and other useful text information to images. The lower half of the menu includes tools for rendering logos and buttons as well as managing the Perl server.

Few of these options are ever used by the average desktop *GIMP* user, though power users will find the DB Browser useful. Additionally, many of the options shown are not part of the stock *GIMP* distribution. They are added by additional plug-ins which can be downloaded from the *GIMP* Registry, various web sites around the world, or acquired from the Graphics Muse Tools CD, a commercial package of over 100 prebuilt plug-ins.

Scanner support

Where to get your drivers

SANE

<http://www.mostang.com/sane/>

XSANE

<http://www.xsane.org/Vuescan>

www.hamrick.com
Image Scan!

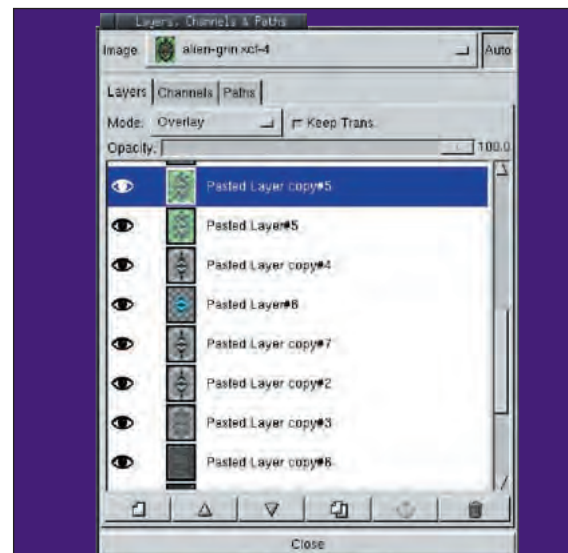
http://www.epkowa.co.jp/english/index_e.html
QuiteInsane

<http://quiteinsane.sourceforge.net/index.shtml>

Epson KOWA

http://www.epkowa.co.jp/english/linux_e/lcd_e.html
Linux USB Project

<http://www.linux-usb.org/>



Managing Layers requires the use of the Layers menu. This menu can be accessed by right-clicking on any layer. Layer management includes four basic functions: creating and deleting, changing the stacking order (including merging layers), managing the alpha channel, and managing the layer mask. A simplified Layers menu is offered in the Canvas menu that will also allow to align layers edges and centers.

Layers Menu

One of the most difficult concepts new users have to conquer is the use of layers. A layer is like a transparent acetate sheet that you paint on. Each sheet is combined in different ways – called 'blend modes' – in order to generate the composited image. Layers have many features and can be moved, resized, and raised and lowered in the stack of layers. The Layers Dialog is used to manage layers and most users will find they have this window opened nearly all the time.




You'll find you use the Layers and Channels dialog often, mostly on the Layers tab. The Channels tab is used by more advanced users to work in a single colour channel to create selections or to save a selection as a new channel. The Paths tab is used to create line drawings that you can later convert into selections or outline using the Edit>Stroke option of the Canvas Menu.

A feature used extensively in large portrait designs is the layer mask. A layer mask blocks out sections of a layer. Light areas of the mask allow the corresponding part of the layer image to show through while dark areas of mask block out corresponding sections. Masks allow you to combine layers in without modifying their contents. This permits a single image, duplicated a number of times to create multiple layers, to be used with different blend modes using different parts of the image without change the original image.

The Canvas Window

The Canvas window is where all your image editing takes place. When you open an existing image, it is displayed as a single layer in a Canvas window. You can edit the first layer – which is called the Background layer by default, or create new layers and edit those.

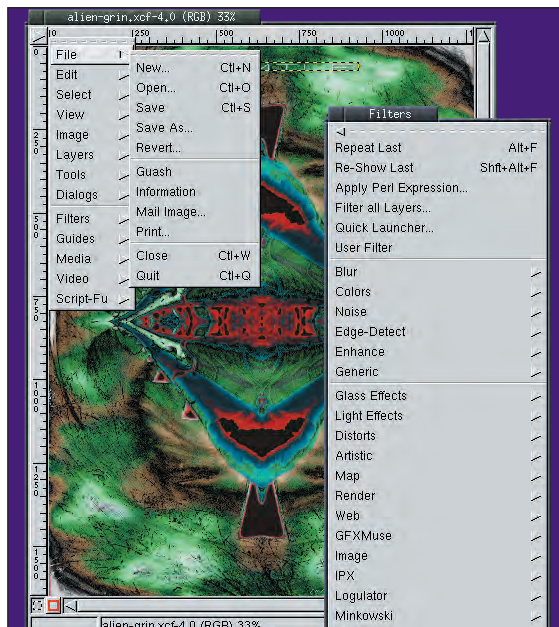
A canvas consists of rulers along the left and top sides of the drawing area, modifiable guide lines that can be pulled from the rulers (hold down the left mouse button and drag from one of the rulers), scrollbars on the right and bottom of the drawing area, a status bar along the bottom of the window, Quick Mask buttons on the lower left and an Image Navigator button on the lower right. The Image Navigator looks like the Move tool icon and allows you to scan around your image if it happens to be larger than the Canvas window currently displays.



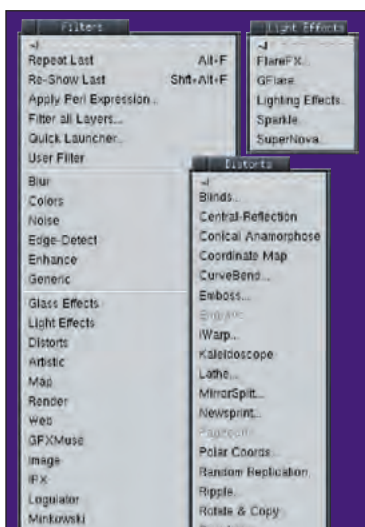
Every canvas has its own Canvas menu. To open this menu, hold down the right mouse button in the image area. The menu offers you a variety of options, from saving the file, to modifying existing selections, to applying filters. This is a tear off menu, which means if you click on the dotted line at

A Canvas window displays the composite image, but it does so based on the settings for the individual layers that comprise the complete image. The Canvas menu is opened by clicking the left mouse button on the arrow button in the upper left corner of the Canvas window or by holding down the right mouse button anywhere inside the Canvas window. When moving an existing Guide Line, be sure to have the Move tool active in the Toolbox. You can instantly resize the Canvas window to fit your image by pressing **Ctrl-E** inside the drawing area of the window.

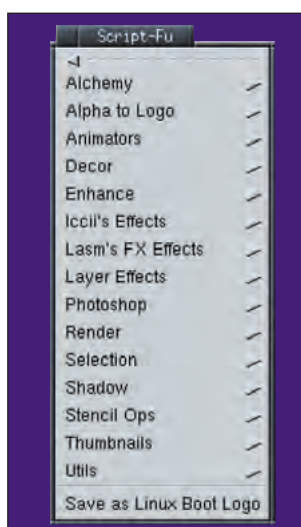
the top of the menu it will become a window of its own that stays open until you click on the dotted line again. All submenus of the Canvas menu are also tear off menus, which means, for example, you could tear off the **Filters>Blur** submenu to have quick access to the multiple blur filters without having to traverse multiple higher level menus.



This menu is used more often than the menus in the Toolbox, so tearing it off to make it a window on its own is probably a good idea. The Script-Fu submenu has an option called "Shadow", which includes an easy to use drop shadow filter. The options shown in the Filters menu here do not all come with the basic **GIMP** installation. Many require you to download, compile and install them manually, or you can purchase prebuilt versions on CD.



There are many filters freely available for **GIMP**, though not all come with the basic distro. The menus make sense for the most part – eg the **Blur** submenu provides filters that blur, the **Distorts** submenu provides filters that twist, flip and otherwise mangle an image and the **Light Effects** render lens flares and white light effects on an image.



Script-Fu is one of a number of scripting languages supported by **The GIMP**. This menu provides a filters which have been written in this language. In essence there isn't any reason to differentiate these filters from filters written in Perl or in any other language. For the short term, just remember this is where you'll find your drop shadow filters!

What's Missing

While **The GIMP** offers many of the same features **Photoshop** users have grown to appreciate, there are still a few missing from this Open Source alternative. The most important of these is proper colour management. The **GIMP** does not currently support true CMYK colour management. In fact, it only works in Grayscale,

Indexed and RGB colour spaces. Additionally, colour channels in **The GIMP** are limited to 8 bits. These two issues combine to make **GIMP** less than optimal for print work. Still, if absolute colour matching is not an issue, **GIMP** can – and does – provide many of the tools necessary for exceptional artwork bound for the Web, in print or in the video industry. **LXF**

Treeware

Books on **The GIMP**

Grokking The GIMP, Carey Bunks, New Riders ISBN: 0-7357-0924-6

GUM – GIMP User Manual, Kylander and Kylander, Coriolis; ISBN: 1-5761-0520-2 (also online at GIMP.org)

GIMP for Web Professionals, Hammel, Prentice Hall ISBN: 0-1301-9114-0

NEXT MONTH

In the coming months we'll be offering a series of tutorials showing how this tool can be used for both the mundane and the sophisticated. In this first tutorial, we'll take a quick run through the windows and dialogs provided to help learn your way around. From here on, we'll assume a basic knowledge of menus and use of the Toolbox and Canvas windows. If you need more detailed information you can check out the online **GIMP** reference guide – the **GUM (GIMP Users Manual)**.

CREATING PDFS

Practical PHP Programming

This month, **Paul Hudson** wraps up the multimedia mini-series by exploring how PHP supports Adobe's ubiquitous Portable Document Format.



The Portable Document Format, or PDF for short, was invented by Adobe as an early attempt to move toward the paperless office – a phenomenon that has still yet to be realised. As a pet project of John Warnock, one of the Adobe founders, PDF received a great deal of backing as a tool that Adobe could use to internally pass documents around. It wasn't until 1991 that "Interchange Postscript", the first public mention of PDF, was shown to the world, then finally released as PDF 1.0 at Comdex Fall in 1992.

That was almost ten years ago, and PDF has come a long way since then – the PDFs of today have fairly good security features, device-independent colour, external links, CMYK colours, JavaScript support, and more. Originally designed to work on all platforms, this ideal has certainly been accomplished – official Acrobat PDF Reader binaries exist for Linux, Windows, Macintosh, AIX, HP-UX, Solaris and OS/2 Warp, and there are a large number of unofficial projects with the ability to read PDFs. There's no shortage of choice!

When it comes to creating PDFs, however, things are more limited. Adobe make a collection of commercial products to the do business, but they invariably come with a hefty price tag and generally are restricted to Windows and Macintosh platforms. Luckily, once again, PHP comes to the rescue!

Before you begin, note that measurements are in *points*, and there are 72 points to an inch. However, this can be altered by changing the output resolution of the produced PDF.

More than one way to do it

Generally speaking, it's fairly rare to find more than one commonly used way to solve a PHP problem involving a specific piece of external functionality. Originally this was true of creating *Flash* movies – one could use either *libswf* or *Ming* to accomplish the same task – and it still is true of creating PDF files.

PHP provides integration with two very distinct PDF libraries – *PDFlib*, by Thomas Merz, and *ClipPDF*, by FastIO. Technically speaking, they are both quite similar, although *ClipPDF* does have the slight edge in that it is multi-threading safe and supports Unicode. Both versions require you to purchase a licence for commercial use, with *ClipPDF* costing \$1000 per IP address and *PDFlib* costing \$500 for a 1-CPU server.

Naturally we don't have the space to cover both libraries here, so for the rest of this article we'll be looking at *PDFlib* – by far the cheaper and more popular of the two.

Creating a PDF document is very much like creating a *Flash* movie or a picture in that to get the desired end result you state the list of drawing actions required to get there – drawing lines, text, adding fonts, etc. As in the prior two articles, you need to track the PDF document you're working with at all times as each function uses it. Let's get started in the usual manner – with a big chunk of PHP!

```
<?php
$pdf = pdf_new();
pdf_open_file($pdf, "/path/to/your.pdf");
$font = pdf_findfont($pdf, "Times-Roman", "host");

pdf_begin_page($pdf, 595, 842);
pdf_setfont($pdf, $font, 30);
pdf_show_xy($pdf, "Printing text is easy", 50, 750);
pdf_end_page($pdf);

pdf_close($pdf);
pdf_delete($pdf);
?>
```

Starting at line one, we use **pdf_new()** to create a new PDF document and store it in **\$pdf**. This value will be used in all the subsequent functions, so it's important to keep.

pdf_open_file() is used to open a file for writing. Note that the free version of *PDFlib* does not allow alteration of existing PDFs; this function merely creates a new PDF of the given filename. Naturally it will need to be somewhere your web server is able to write to, otherwise you'll receive an error along the lines of "Fatal error: PDFlib error: function 'PDF_set_info' must not be called in 'object' scope in script.php on line XYZ"

The next line uses **pdf_findfont()** to find and load a font for use inside the generated PDF file. In the example on this page, **pdf_findfont()** takes three parameters – the PDF document to work with, the name of the font to use, and which encoding to use. In the example above, **\$pdf** is specified as the first parameter (as always). "Times-Roman" is specified as the font to use, which is one of the fourteen standard internal *PDFlib* fonts. The next parameter can be set to either **winansi** (Windows), **macroman** (Macintosh), **ebcdic** (EBCDIC code page 1047 machines), **builtin** (for symbol fonts), or **host** (winansi for Windows, macroman for Macintosh, etc). Generally it's smart to use **host**, and leave *PDFlib* to do the hard work, so this is what is used in the code above.

Special note

From this issue on, I'm using PHP 4.3.0 for development and testing. If you find something doesn't work on your system, try upgrading to a newer release.

pdf_findfont(), when successful, returns a font resource which is stored in **\$font**. You may wish to add error-checking in your own scripts to add extra reliability.

At this point, we're ready to start on the main part of PDF generation. The first three lines merely set things up for the document – the next four, lines four to seven, are the page itself.

Reading the source, it's fairly easy to see that line four and line seven encapsulate one page in the generated PDF file. Objects and text outputted between a **pdf_begin_page()** and **pdf_end_page()** will affect that page, and multiple begin/end blocks are used to create multiple pages. Note that **pdf_begin_page()** takes a second and third parameter which specify the X and Y point size of this page. The PDF format allows you to make your pages different point sizes from page to page – read the *Point sizes in real life* box on the right for some commonly used sizes.

pdf_setfont() takes three parameters – the first is as per usual, the second parameter is the return value from **pdf_findfont()** for the font you wish to use, and the final parameter is the size to use, in points. Immediately afterwards, we call **pdf_show_xy()** to place text into our page. Parameter two of **pdf_show_xy()** is the string to use, and parameters three and four are the X and Y co-ordinates at which to print the text. It's important to note that the Y parameter, the last parameter passed to **pdf_show_xy()**, is distance the text should appear above the page baseline, in points. That is, setting this parameter to 0 will have the bottom of a lowercase "a" at the very bottom of the page, and the bottom of a lowercase "y" outside the margins of the page.

With **pdf_end_page** called, the first and only page is completed, and all that's left to do is clean things up.

Cleaning up is done through the help of two functions – **pdf_close()** and **pdf_delete()**. They may sound somewhat similar, but you do need to call them both – **pdf_close()** cleans up the *PDFlib* memory and document-related resources, whereas **pdf_delete()** cleans up PHP's reference to **\$pdf** and any other internal resources. Be sure to call them in the order shown above.

Save the code as *lxpdf1.php* and, after placing inside your public HTML directory, run it using your web browser. Note that there's no "Success!" message being printed out, however, you should find your PDF file has been created and is viewable in your PDF reader of choice.

Adding more pages and more style

At this point, our script enables us merely to display one page – here's a simple modification to demonstrate multiple pages:

```
for ($i = 1; $i < 10; ++$i) {
```

Installing PDF support

The easiest part of generating PDF files...

First, grab the JPEG, TIFF, and, optionally, the PNG libraries. Secondly, download and install *PDFlib* from <http://www.pdflib.com/pdflib/>. Finally, add the following four options to your */configure/* line:

```
--with-pdflib=/path/to/pdflib
--with-jpeg-dir=/path/to/libjpeg
--with-png-dir=/path/to/libpng
--with-tiff-dir=/path/to/libtiff
```

If you followed the *LXF* installation instructions back in issue 30, you will already have *PDFlib* installed and so would only need to repeat the above if you wish to change other settings or upgrade.

This is page 1

Subheader

This is some standard-sized text in Times-Italic

```
pdf_begin_page($pdf, 595, 842);
pdf_setfont($pdf, $font, 30);
pdf_show_xy($pdf, "This is page $i", 50, 750);
pdf_end_page($pdf);
}
```

As you can see, it is simply a matter of calling **pdf_begin_page()** and **pdf_end_page()** repeatedly – although you will likely want to make the content inbetween more interesting!

A good start is to have a selection of type faces ready for various parts of your document. In our first example, we have just one – Times-Roman is stored in **\$font**. However, that could be easily modified as such:

```
$times = pdf_findfont($pdf, "Times-Roman", "host");
$timesb = pdf_findfont($pdf, "Times-Bold", "host");
$timesi = pdf_findfont($pdf, "Times-Italic", "host");
```

Combined with the use of **pdf_setfont()**'s third parameter, we can therefore create headers and subheaders like this:

```
for ($i = 1; $i < 10; ++$i) {
pdf_begin_page($pdf, 595, 842);
pdf_setfont($pdf, $times, 24);
pdf_show_xy($pdf, "This is page $i", 50, 750);
pdf_setfont($pdf, $timesb, 16);
pdf_show_xy($pdf, "Subheader", 100, 700);
pdf_setfont($pdf, $timesi, 16);
pdf_show_xy($pdf, "This is some standard text", 100, 700);
pdf_end_page($pdf);
}
```

We can even throw in the **pdf_setrgbcolor()** function, which takes a red, green, and blue value as its second, third, and fourth parameters respectively, and uses them to set the colour of fills and objects that follow.

Try adding this line just before the first **pdf_setfont()**...

```
pdf_setrgbcolor($pdf, 1.0 - (0.1 * $i), 0.0, 0.0);
```

And adding this line just before the second **pdf_setfont()**...

```
pdf_setrgbcolor($pdf, 0.0, 0.0, 0.0 + (0.1 * $i));
```

Save the file as *lxpdf2.php* and load it up in your browser – your output should be very close to the picture at the top of this column – if not, you might want to check your code for errors.

All being well, your top header should start off red and fade into black, whereas your second-level header and the main text should start off black and fade into blue.

Adding imagery

Let's face it – there's only so far you can go with plain text, no matter *how* many font faces and colours you use. Indeed, generally you should try to limit typeface diversity if possible, because very often adding more words won't improve the presentation of your work – which is perhaps why the received wisdom is that a picture says a thousand words. Even if the words are multicoloured and in a wacky font.

To this end, PHP provides us with two functions:

pdf_open_image_file() and **pdf_place_image()**. The former reads a specified image type (parameter two) of a specified file name (parameter three) and returns an image that can be used in subsequent functions.

Headers, subheaders, and colours all at once. Best not to overdo it, though – less is more!

Point sizes in real life

Here's a short list of popular paper sizes, in points

A0	2380x3368
A1	1684x2380
A2	1190x1684
A3	842x1190
A4	595x842
A5	421x595
A6	297x421

Remember, the PDF standard is marvellously flexible in that it allows you to have very varied page sizes; don't feel constrained to stick with A4 all the time!

PDF without a module?

If you're interested in a unique way to create PDFs, point your web browser at [pdf-php](http://pdf-php.sourceforge.net/projects/pdf-php) – <http://sourceforge.net/projects/pdf-php> – the project allows you to create PDF documents entirely using PHP.



« **pdf_place_image()** then takes the returned image as its second parameter, and also allows you to specify the X co-ordinate (parameter three), Y co-ordinate (parameter four), and any scaling (parameter five) you wish to be applied to the image.

Save this next example as `lxfpdf3.php`. You will need to find a JPEG and place it in the same directory as the script as `myimage.jpg` before you run the script.

```
<?php
$pdf = pdf_new();
pdf_open_file($pdf, "/path/to/your.pdf");

pdf_begin_page($pdf, 595, 842);
$testimage = pdf_open_image_file($pdf, "jpeg",
    "myimage.jpg");
pdf_place_image($pdf, $testimage, 0, 0, 0.5);
pdf_end_page($pdf);

pdf_close($pdf);
pdf_delete($pdf);
?>
```

In the above example, we set the *scale* parameter of **pdf_place_image()** (parameter five) to 0.5, which will show our `myimage.jpg` picture at 50% of its usual size. Note that altering the scale value of pictures *will not* change the final file size of the PDF that you output, because the file is saved *unscaled* then scaled at run-time.

This might seem odd, but it makes good sense when you realise the reason behind it. Owing to its saving pictures unscaled, the PDF format allows you to re-use images without having to store multiple copies in the file. So, if we go back to our earlier **for** loop where we had ten pages being generated, we get something like this:

```
<?php
$pdf = pdf_new();
pdf_open_file($pdf, "/path/to/your.pdf");
```

PDF Resources

Interested in learning more?

As we often find ourselves saying at *LXF*, four pages is rarely enough to get across full coverage of a subject, and PDF generation is no exception. Due to its relative complexity, and the fact that at the end of the day it is still a fairly niche solution, there are not all that many PHP/PDF documents online.

However, of particular interest when it comes to online articles are the following:

Perugini Luca has written a quick but helpful introduction to PHP/PDF/MySQL online at <http://www.phpbuilder.com/columns/perugini20001026.php3>

Paul Adams has written an equally quick introduction to PDF creation with PHP online at <http://hotwired.lycos.com/webmonkey/02/20/index3a.html>

If you're struggling to expand your knowledge beyond basic primers like these and don't fancy shelling out for one of the multitude of books on the PDF format, you might have to head over

to the Adobe site to read the official documentation. You can grab the PDF standard, a PDF document itself being used for exactly its originally intended purpose from

<http://partners.adobe.com/asn/developer/acrosdk/docs/fileformats/PDFReference.pdf> – but beware, it's quite large.

Finally, perhaps your best bet to learn as much about PDF with PHP as you can is the PHP manual itself. UK readers can take a look at <http://uk.php.net/manual/en/ref.pdf.php> – if you're from elsewhere you will probably want to use your local mirror.

With regards to books that give good coverage of PDF/PHP, *Mastering PHP 4.1* (reviewed last issue, written by Allen and Hornberger and published by Sybex ISBN 0-7821-2924-2) has about forty pages dedicated to PDF generation, and is an excellent choice for this topic. Peter Moulding's *PHP Black Book* (Coriolis, ISBN 1-5888-0053-9), in comparison, is fairly limited in its coverage of PDF creation – although it does cover *ClipPDF* to a good extent.

ClipPDF?

Interoperability

Happily, **PDFlib** and **ClipPDF** aren't all that far apart when it comes to using the two inside PHP: once you're familiar with **PDFlib**, it will probably be worth your while investigating whether **ClipPDF** is better suited to your needs. One minor drawback to using **ClipPDF** is that the PHP documentation is

a great deal less complete for **ClipPDF** than it is for **PDFlib**, which would make learning it much trickier. Have a look over the documentation at <http://uk.php.net/manual/en/ref.cpdf.php> (or your local mirror) and see what you think.

```
$times = pdf_findfont($pdf, "Times-Roman", "host");
$timesb = pdf_findfont($pdf, "Times-Bold", "host");
$timesi = pdf_findfont($pdf, "Times-Italic", "host");
$testimage = pdf_open_image_file($pdf, "jpeg",
    "myimage.jpg");
```

```
for ($i = 1; $i < 10; ++$i) {
    pdf_begin_page($pdf, 595, 842);
    pdf_setrgbcolor($pdf, 0.0, 0.0, 0.0);
    pdf_setfont($pdf, $times, 24);
    $scaleval = $i * 10 . '%';
    $smallscale = 0.1 * $i;

    pdf_show_xy($pdf, "This is page $i - $scaleval scale", 50,
        750);
    pdf_place_image($pdf, $testimage, 0, 0, $smallscale);
    pdf_end_page($pdf);
}

pdf_close($pdf);
pdf_delete($pdf);
?>
```

Save that as `lxfpdf4.php` and run it again – this time the `your.pdf` file should be only slightly larger than that generated by the previous script that had just one image. As you can see, it's actually quite helpful that PDF stores its graphics separately from the display, particularly when you re-use pictures a great deal.

We can further manipulate images through the use of **pdf_rotate()** and **pdf_skew()** – two functions you should be able to guess quite easily.

As per usual, both **pdf_rotate()** and **pdf_skew()** take a PDF document reference as their first parameter. **pdf_rotate()** then takes one extra parameter – how much to rotate the co-ordinate system by, in degrees – whereas **pdf_skew** takes two extra parameters – how much to skew the co-ordinate system by in the X direction and how much in the Y direction.

Try adding these two lines below just after the call to **pdf_begin_page()** inside the loop:

```
pdf_skew($pdf, 10, 10);
pdf_rotate($pdf, 5);
```

Re-save the file and generate the PDF file again. You may need to comment one of the lines out at a time to clearly see what effect each call has on your document, however you should also notice that these two functions allow you to rotate and skew your text as well as your pictures, which can add much extra visual impact to your PDFs.

Adding document data

Owing to the fact that PDFs are designed to be read like normal, printed documents, Adobe have incorporated into the PDF format the ability to add notes in the same manner one might scribble in a margin.

These notes, which can be edited and re-edited by readers, can also be created using PHP by calling the function

pdf_add_note(). Here's an example of use:

```
pdf_add_note($pdf, 100, 500, 700, 600, "You can create
notes very easily using pdf_add_note()", "Sticky notes", "note",
1);
```

The second, third, fourth, and fifth parameters are respectively the lower-left X and lower-left Y co-ordinates, and the upper-right X and upper-right Y co-ordinates of the note boundaries. The sixth and seventh parameters are the text to put inside the note and the title to place at the top, and the final two parameters decide the icon used to display the note when closed and whether or not the note starts open. Note that once the PDF is loaded, your reader is often free to move these notes around and edit the text inside them.

So, in the line above, we add a 600x100 note box which is already open (use 1 to specify the note is open, and 0 to specify it is closed). Instead of **note** as the penultimate parameter, we have various other options: **comment**, **insert**, **paragraph**, **newparagraph**, **key**, or **help**. In several PDF readers, this parameter has no effect and can be just left as **note**.

Another important facet to improving the usefulness of documents is to provide meta-data regarding who created the document and when. This can be achieved through the use of **pdf_set_info()**, which takes a key and a value as its second and third parameters. The standard keys for use as **Subject**, **Title**, **Creator**, **Author**, and **Keywords**, but you are also able to add your own keys such as **Modified**, **Created**, etc.

Now we can finish off our script by adding in some meta-data – add these three lines just below **pdf_open_file()** and save the modified version as **lxfpdf5.php**:

```
pdf_set_info($pdf, "Creator", "Hudzilla");
pdf_set_info($pdf, "Title", "PHP PDF 101");
pdf_set_info($pdf, "MyInfo", "You can write what you please
here");
```

When you read the PDF generated by **lxfpdf5.php** you should see the note sticking out in quite an obvious fashion. The meta-data will be there too, but it is likely to be hidden away under a menu somewhere.

Conclusion

In this PDF primer we've explored how to create documents and pages inside them, add fonts of various sizes and colours, add and manipulate images, and finally how to add non-paginated information to the document. There is a fair amount more to working with PDF files than we've been able to cover here, however the core concepts are all above. There is an excellent forum at <http://groups.yahoo.com/group/pdflib> where you can ask further questions about using *PDFlib*, and it's well worth a look if you intend on delving into more advanced areas.

Adding dynamically generated PDFs to your site can be a great boon, especially where official documents are concerned – it's definitely worth the pay-off to take the time to get PDF generation down to a fine art in order to add another valuable item to your PHP toolkit. Admittedly, there are still some



Skewing images is easy, as you can see.

limitations in the technology: the inability to load existing documents for editing is the one most commonly cited. However, that disadvantage is almost certainly outweighed by the sheer popularity of the format. With this, we come to the end of the multimedia mini-series. I hope you've learnt a lot about how much variety you can add to your pages through the use of PHP, and, perhaps more importantly, how little the generation of images, Flash movies, and PDF document differs behind the scenes. I'd be interested to hear whether any readers fancied creating a unified library in PHP where text and shapes can be manipulated at a higher level than images, *Flash*, and PDF, and where the output format was only decided at the end of the script. Over to you...! Good luck! [LXF](#)

About Paul Hudson

Paul Hudson is a London-based web developer specialising in PHP and Perl. He can be emailed at hudzilla@php.net

NEXT MONTH

In issue 39, I'll be going into the unknown – the only thing I'll say is that the topic will be "Alternative uses for PHP". Don't let your imagination get carried away – but certainly be first in the queue at your local newsagents! :-)

If you have any comments or suggestions about this series, please be sure to write in.



PYTHON PASSION

Crazy about Python

Patrick K. O'Brien launches our Python tutorial series by confessing his great affection for the language.

Notable Quote

From the creator himself

"The joy of coding Python should be in seeing short, concise, readable classes that express a lot of action in a small amount of clear code – not in reams of trivial code that bores the reader to death."

Guido van Rossum, creator of Python

Let's get one thing out in the open: I love Python. I realise you may think love is too strong a word to describe the relationship between a programmer and his programming language. If that's the case, I intend to prove you wrong. In fact, over the next several months I intend to show you all the ways in which Python is worthy of my love. Why am I doing this? I'll tell you why: I want you to love Python, too.

Before you get the impression that I'm a hopeless romantic who's oblivious to the flaws of his obsession, let me assure you that isn't the case. I've had my share of relationships: Cobol, RPG, dBase, Paradox, C, C++, Delphi, FoxPro, Clarion, Perl (a one-night stand that ended badly for both of us), Visual Basic, and Java. I've had my crushes, infatuations, and passionate moments. I've also been spurned, rejected, and stood up. At various points I thought about giving up programming altogether: it was simply too painful. Then I found Python.

Python was everything I had been looking for in a programming language: lightweight, versatile, elegant, and

undemanding. Better still, Python had all the qualities of the best Open Source software: free as in beer, and free as in freedom; stable and reliable beyond any programming language I'd ever used before; portable across all the popular operating systems, such as Linux, Windows, and Mac; and enriched by a wonderfully supportive community of programmers. I was charmed and fascinated; over time, my infatuation matured into admiration.

While that admiration continues unabated, my attraction to Python isn't blind. I recognise that Python isn't perfect; she has a blemish or two, and a few scars from childhood. After all, Python was created by Guido van Rossum around 1990, making Python older than both Linux and Java, and nearly as old as Perl. But Python has aged well, and recent surgery (version 2.2's type/class unification, for example) has brought new vitality and even corrected some troublesome birth defects. Besides, beyond superficial attractions, shouldn't one look to personality traits in choosing a companion? After all, we are talking about a relationship that could extend for years. Surely those should be happy, productive years. In my own case, reflecting on the time I've spent with Python, the most significant thing I can say is that while Python and I have sometimes disagreed, we've never quarrelled, and unlike my previous relationships, I've never raised

Recommended Resources

Places to go, things to see, people to meet

Official Python Website:

<http://www.python.org/>

Official Python Documentation:

<http://www.python.org/doc/>

Python Newsgroup:

<news:comp.lang.python>

Python Mailing Lists:

<http://www.python.org/psa/MailingLists.html>

Python Tutor Mailing List:

<http://mail.python.org/mailman/listinfo/tutor/>

Patrick K. O'Brien Website:

<http://www.orbtech.com/web/pobrien/>

my voice in anger or frustration. Instead, Python has renewed my enthusiasm for programming, and put fun back into my work.

Getting Acquainted

How could you not love a language that's powerful and versatile, yet makes programming fun? One that gives you a feeling of power and control? One that satisfies your desire for beauty and elegance, without letting you down in a pinch? I couldn't. And that's why I say that I love Python. Because, for me, Python brought back the feelings I had when I first started working with computers some 15 years ago: the feeling that I could design and create tools that did incredible things, things that others found amazing and almost magical. That's a wonderful feeling, and it's one I'd like to share with you.

To help you join in that feeling, I'm going to introduce you to Python through a series of tutorials. Even better, I'll act as chaperone while you and Python get to know each other: I'll point out Python's virtues, give you advice for making the most of your relationship, and highlight other resources that might be of interest to you.

Whether you've programmed before, or have only thought about programming, you'll find Python to be an approachable language. The syntax is easy to learn, easy to write, and easy to read. (You'll most appreciate this quality when you revisit code that you wrote months ago.) Python's rules are simple and logical, with very few exceptions to trip you up; Python rarely behaves in ways that would be considered unusual or unexpected. And its constructs are concise, powerful and versatile. (Once you understand lists and dictionaries, for example, you'll use them as often as a carpenter uses a hammer and nails.) As you master the basics, you'll begin to see that you can use Python to create a wide variety of applications, from simple utilities to complex, mission-critical systems.

Let's begin this courtship, then, by looking at an example program that will allow you to form a first impression of Python.

Article Setup

I find that having a personal problem to solve is the best motivator for writing a program. The most recent problem I needed to solve involved this very magazine. After I agreed to write a series of Python tutorials, and the editor had approved the topics, I then had to set about writing them. And I couldn't

just write them any which way; I had to submit text files that conformed to the tagging conventions this magazine uses to denote elements such as headings and body copy, code listings and screen grabs. As a programmer at heart, I do my best to automate many of my tasks, even article writing. So I wrote a Python program (named `articlesetup.py`) to create a directory and initial text file for each of the tutorial topics I was going to cover. While I was at it, I also had the program create a `.cvsignore` file for each directory. The result was this:

```
# The os module provides common operating system services.
# import os topics is a list (an ordered sequence) of article
topics = ['Intro', 'Basics', 'Containers', 'Flow',
          'Packaging', 'Functions', 'Classes']

# template is text I'd like to appear in all my articles.
# These are all the basic tags required by the magazine.
# Python triple-quoted strings retain embedded newlines.
template = """
///DESCRIPTION///
†
///TITLE///

///STRAP///

///BODY COPY START///

///CROSSHEAD///

///END BODY COPY///

///NEXT MONTH///
"""

# This is the main function defined by this program.
# It is common to name it 'main', but you don't have to.

def main():
    for topic in topics:
        # Dirs should look like '1.Intro', '2.Basics', etc.
```

Notable Quote

From an Open Source legend

"These days I write Python when I can and C when I must. Despite minor flaws, I think Python is the most expressive and powerful language I have ever used; it has even seduced me away from my old-time loyalty to LISP." — Eric S. Raymond, Open Source evangelist and author of "The Cathedral and the Bazaar"

www.python.org should be bookmarked by anyone serious about using the language.



TutorialPython

```
[pobrien@localhost pobrien]$ python
Python 2.2.2 (#4, Dec 13 2002, 21:58:45)
[GCC 3.2 (Mandrake Linux 9.0 3.2-1mdk)] on linux2
Type "help", "copyright", "credits" or "license" for more information.
>>> 2 + 2
4
>>> x = 3
>>> y = 7
>>> x * y
21
>>> greeting = 'Hello.'
>>> question = 'How are you?'
>>> people = ['Susan', 'Bob', 'Mary', 'Steve']
>>> people.sort()
>>> people
['Bob', 'Mary', 'Steve', 'Susan']
>>> for person in people:
...     print greeting, person + '!', question
...
Hello, Bob! How are you?
Hello, Mary! How are you?
Hello, Steve! How are you?
Hello, Susan! How are you?
>>>
```

Running Python in interactive mode is an excellent learning aid, as it lets you communicate with the interpreter and get immediate feedback every step of the way.



```
# List indexes start at zero, so we add one.
dir = str(topics.index(topic) + 1) + ' ' + topic
if not os.path.exists(dir):
    os.mkdir(dir)
# Tell CVS to ignore compiled Python files: *.pyc.
filename = os.path.join(dir, '.cvsignore')
if not os.path.exists(filename):
    f = file(filename, 'w')
    f.write(".pyc\n")
    f.close()
# Article files should look like 'Intro.txt', etc.
filename = os.path.join(dir, topic + '.txt')
if not os.path.exists(filename):
    f = file(filename, 'w')
    f.write(template)
    f.close()

# The following `if` statement is a common idiom in Python.
# If this module is sent to the Python interpreter as a
# command line argument, its `__name__` attribute will be
```

```
articlesetup.py
# The os module provides common operating system services
import os

# topics is a list (an ordered sequence) of article topics
topics = ['Intro', 'Basics', 'Containers', 'Flow',
          'Packaging', 'Functions', 'Classes']

# template is text I'd like to appear in all my articles
# These are all the basic tags required by the magazine
# Python triple-quoted strings retain embedded newlines
template = """
//DESCRIPTION//

//TITLE//

//STRAP//

//BODY COPY START//

//CROSSHEAD//

//END BODY COPY//

//NEXT MONTH//
"""

# This is the main function defined by this program.
# It is common to name it 'main', but you don't have to.
def main():
    for topic in topics:
        # Dir should look like '1.Intro', '2.Basics', etc.
        # List indexes start at zero, so we add one.
        dir = str(topics.index(topic) + 1) + ' ' + topic
        if not os.path.exists(dir):
            os.mkdir(dir)
        # Tell CVS to ignore compiled Python files: *.pyc.
        filename = os.path.join(dir, '.cvsignore')
        if not os.path.exists(filename):
            f = file(filename, 'w')
            f.write(".pyc\n")
            f.close()
```

I wrote `articlesetup.py` in XEmacs to make sure my articles are tagged to the LXF Production Editor's specification.

```
# set to `__main__`, and the `main()` function will be
# called. If, instead, this module is imported by another
# program, its `__name__` will be 'articlesetup' and the
# 'main()' function will not be called.
```

```
if __name__ == '__main__':
    main()
```

This program represents the sort of simple utility that I would normally write, run once, and delete. But since I wanted to include it in this article, I added comments (those lines beginning with a `#`) that I wouldn't normally have bothered with in this kind of program. I also made the code slightly more sophisticated than it needed to be, in order to illustrate some features of the Python language. One of those features is support for several different programming paradigms, such as procedural, functional, and object-oriented. Don't worry if those terms are unfamiliar. For now we'll just examine this one program to see what it reveals about Python.

Priceless Projects

Gadgets, contraptions, and other treasures

wxPython – www.wxpython.org/

If you want to create graphical user interfaces that run on Linux, Windows, and the Macintosh, using native controls on each, wxPython is hard to beat. Even if you don't, download it and install it anyway because it includes *PyCrust*, which we will use throughout this tutorial series.

Zope – www.zope.org/

For those of you interested in web application development, Zope is the 800-pound gorilla of content management systems. Zope comes with a web server, object database, templating system, relational database support, and more. Those who master it are said to have acquired the Zen of Zope.

Quixote – www.mems-exchange.org/software/quixote/

If a one-ounce pygmy mouse lemur is more your style, you'll find Quixote to be a practical alternative to Zope for building web interfaces to your Python programs.

Twisted – www.twistedmatrix.com/

Asynchronous application development is incredibly difficult to do right, unless, of course, you use the Twisted framework. With built-in support for just about every network protocol, you can use Twisted to create both clients and servers for the web, email, ftp, Internet Relay Chat, Instant Messaging, remote object access, and more.

Program files and module files

Let's start with the program file itself. You may have noticed that the code refers to itself as a module, while in this article I've been calling it a program. In Python, the distinction between a program and a module is trivial, and for a good reason: any file that contains Python code is a module and may be used by other modules. One module uses code contained in another by importing it, such as I did when I imported the `os` module. The only thing that makes a module also a program is the fact that it contains code that can be run from the command line. In our case, `articlesetup.py` meets that criteria.

I created the `articlesetup.py` file using a regular text editor (XEmacs, in my case) and saved the file to the directory in which I wanted the program to create new directories. Then I ran it from the command line. To run a Python program from the command line you simply pass the program filename as an argument to the Python interpreter:

```
$ python articlesetup.py
```

Notable Quotes

Python Success Stories – www.pythonology.com/success

"Python plays a key role in our production pipeline. Without it a project the size of Star Wars: Episode II would have been very difficult to pull off. From crowd rendering to batch processing to compositing, Python binds all things together." – *Tommy Burnette, Senior Technical Director, Industrial Light & Magic*

"We chose Python because it provides maximum productivity, code that's clear and easy to maintain, strong and extensive (and growing!) libraries, and excellent capabilities for integration with other applications on any platform." – *Steve Waterbury, Software Group Leader, NASA STEP Testbed*

"Our project was a huge success. We rebuilt eight years

of software development effort from the ground up in less than two years, and with a substantially smaller team. I attribute this largely to Python. With Python, it is very easy to develop code quickly." – *Michael Muller, Philips*

"Python has been an important part of Google since the beginning, and remains so as the system grows and evolves. Today dozens of Google engineers use Python, and we're looking for more people with skills in this language." – *Peter Norvig, Director of Search Quality, Google, Inc.*

"Python made it possible for Journyx to produce a flexible, feature-rich product for multiple platforms in

less time than would have been possible using another language. The Timesheet project has succeeded spectacularly, generating millions in revenue and allowing Journyx to grow every year, even under the current economic conditions." – *Curt Finch, Founder and CEO, Journyx*

"HomeGain maintains its commitment to continual improvement through rapid turnaround of new features and enhancements. Python supports this short time-to-market philosophy with concise, clear syntax and a powerful standard library. New development proceeds rapidly, and maintenance of existing code is straightforward and fast." – *Geoff Gerriets, Software Engineer, HomeGain.com*

When I sent the `articlesetup` module to Python as a command line argument, its `__name__` attribute was set to `__main__` by the Python interpreter. That resulted in a positive equality test in the `if` statement at the bottom of the program, which then invoked the `main()` function. The `main()` function iterated over the list of topics, creating a directory, a `.cvsignore` file, and an article text file for each topic. The article text files were primed with text defined in the string variable named `'template'`. After the program ran, I was ready to begin writing the article you are now reading.

Program Execution

You probably noticed that my description of the running program began with code located at the bottom of the program, and worked its way back to the top. That's because the Python interpreter reads a program file the same way we do: line by line, starting with the first line in the file and working down to the last line. That means variables and functions must be defined earlier than the code that makes use of them. It also means that you typically have to scroll to the bottom of a program file to find its starting point.

For example, it is quite common for a module to serve two distinct purposes: as a program run from the command line, and as a resource to be used by other modules. Modules that are designed for both purposes switch between these two modes of execution by performing a simple test, which you will frequently find at the bottom of a Python program:

```
if __name__ == '__main__':
    main()
```

Are you curious about the rather strange looking `__name__` attribute? Python distinguishes special, reserved attributes by starting and ending them with double underscores, as a way of keeping them from conflicting with attributes that you might create yourself. Ordinary user-defined variables and attributes can have ordinary-looking names: I named my list of topics just that – **topics**; the text that I wanted to appear in each text file was named **template**; and I called the primary block of functionality **main**. I took the same approach with my **dir** and **filename** variables. And creating these variable was as simple as assigning a value to the variable name.

Consider what I didn't have to do to make use of these variables: I didn't have to tell Python what type of information they would contain, or allocate memory for their contents, nor did I have to clean up after I was finished with them. Other

languages require you to worry about these aspects of a running program, and getting it wrong means your programs will crash or behave badly in ways that are often difficult to debug. Python takes care of those details with its dynamic typing, automatic memory management, and garbage collection.

Now we've experienced writing and executing a Python program, let's look at the other main way to use Python: interactively.

Interactive Mode

A quick way to check whether you have access to the Python interpreter is to enter **python** at the command line. If you do have access, you'll see something like this:

```
$ python
Python 2.2.2 (#4, Dec 13 2002, 21:58:45)
[GCC 3.2 (Mandrake Linux 9.0 3.2-1mdk)] on linux2
Type "help", "copyright", "credits" or "license" for more info.
>>>
```

What we've just done is start Python in interactive mode. Python responds by displaying a few lines of text containing its version, the operating system, and ways to get additional information, followed by a prompt (`>>>`). You can now enter a line of Python code and hit **Enter**, at which point Python will evaluate the code and display either a response or a secondary prompt (`...`) if the code you entered was the start of a block that requires additional lines of code (such as an **if** statement or **for** loop). Python's interactive mode is also known as the Python shell. To exit the Python shell type **Control-D** at the primary prompt.

Running Python in interactive mode is a great way to learn Python. You can use the Python shell as a simple calculator to evaluate math expressions, as a scratch pad to try out new code, or as a tool to explore the functionality provided by Python's standard library of modules. I'm such a big fan of Python's interactive capabilities that I actually wrote my own graphical Python shell, called *PyCrust*. So it shouldn't come as a surprise that I plan to focus on the Python shell throughout this tutorial series. But I'll go one step further and predict that you won't be completely won over by Python until you experience her interactive aspects. After all, communication is the foundation for every enduring relationship, and Python is a great conversationalist.

Ready For Romance?

Here is my parting advice to you: go ahead, take the plunge, fall in love... with Python. [LXF](#)

About the author

Python expert Patrick O'Brien is the author of *PyCrust*, an interactive, graphical Python shell, a developer of the *PythonCard wxPython* application construction kit, and leader of the *PyPerSyst* team, which is working on object persistence mechanisms for Python applications. His articles have been published by O'Reilly and IBM, and he has presented at the O'Reilly Open Source Convention. You can contact Patrick at pobrien@orbtech.com

NEXT MONTH

We'll cover the basics of Python in detail: installation; interactive mode; basic data types, operators, and expressions; and importing modules from the standard library. We'll also introduce *PyCrust*, the flakiest Python shell.

BACKUPS MADE EASY

Backing up partitions

Backing up isn't hard to do with **Nick Veitch's** guide to partition imaging, so go and save your data!

We have often gone on at great length about the necessity of keeping backups as protection against unforeseen events. You don't realise how valuable your data is until you miss it!

Backing up files is one thing, but it is often more useful to back up entire partitions. Restoring the root partition of a drive from a file backup can be tricky, and what do you do if the host system has no viable installed OS?

The solution is to back up entire partition images, although file backups are also useful. There are further advantages to this approach. If you have a set of identical client machines (in a web café environment, for example), you can simply have one backup image which can be used to restore any of the clients. This can also make upgrading clients much easier – install and configure the software correctly on one and you can then use an image to update all the others. With the right software, you could even do this over a network, simultaneously.

Using the dd tool

For simple partition backups, you might have used the *dd* tool. This versatile tool is used for copying data blocks from one device to another or to a file. For example, to back up your /boot partition which lives on /dev/hda1 you might use:

```
dd if=/dev/hda1 of=$HOME/boot.img
```

This quite simply copies the blocks of /dev/hda1 until it reaches the end. It is possible to specify further options for *dd* of course, but these often aren't useful for copying partitions.

While *dd* is perfectly useful for small images like floppy disks, boot disks and your /boot partition, it isn't quite so useful when it comes to larger partitions. For example, if you want to duplicate a 2GB partition you are going to need another 2GB of filespace somewhere else. This is still the case even if only a small portion of the partition is actually in use. You can retroactively compress the image file, but you still need the disk space to begin with.

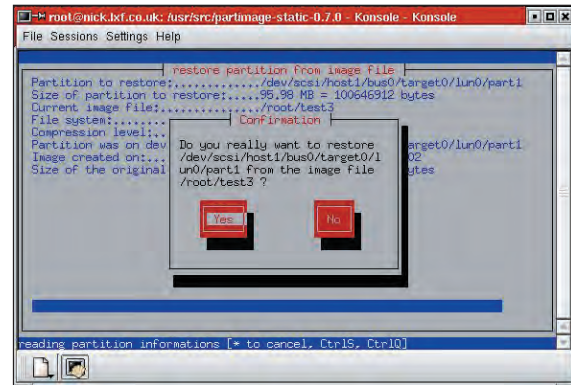
Enter Partition Image

These are the kind of limitations that *Partition Image* was developed to overcome. It can dramatically reduce the time and space required to complete partition backups in two ways:

- It only backs up used blocks in the partition. This alone is usually a major space saver. Few people run their partitions to 100% full, so there is space to be saved here.

- The software can also use a variety of compression tools to reduce the size of the image being created.

For all but the very simplest of image files, using *Partition Image*



Restoring a partition should be straightforward, and quicker than the original backup.

Required libraries

libz	gzip compression library www.zlib.org
libbz2	bzip2 compression http://sources.redhat.com/bzip2/
liblz	fast compression library http://www.oberhumer.com/opensource/lzo/
libnewt	curses-like GUI system ftp://ftp.redhat.com/pub/redhat/linux/code/newt/
libslang	slang, required for newt http://www.s-lang.org
libssl	OpenSSL http://www.openssl.org
libmcrypt	Mcrypt for secure networking http://mcrypt.hellug.gr

is much more effective, both from a time and storage point of view. The time penalty incurred by compressing the data is often made up for by the fact that you are compressing a smaller file, and the subsequent reduced time taken to copy the image onto backup devices where necessary.

The story so far

Francois Dupoux and Franck Ladurelle released version 0.2 of *Partition Image* in May of 2000. Since then it has grown from its fairly basic beginnings to support more and more filesystems, additional compression formats, and an easy to use GUI, based on the curses-like *newt* system.

You can download *PI* from the homepage, which is found at www.partitionimage.org. There are a number of ways to install *PI*, and a number of packages available. If you just want to run the latest stable version, you are probably best off downloading the static binaries.

If you want to compile it yourself either because there isn't a binary available or you want the latest CVS version, there are a few libraries you will need – see the *Required Libraries* box.

To compile from source (once have installed the libraries), you can follow the usual procedure of running **configure** and **make**. However, you may want to configure with options for additional filesystem support, or to control the location of SSL or disable SSL on the server. Running **./configure --help** will give you a full list of the options. So, for example, you might want to say

```
./configure --enable-xfs
```

```
make
```

```
make install
```

If you have enabled SSL, you'll want to make the certificates:

```
make certificates
```

Partition tables

Before we continue to talk about backing up partitions, we need to discuss how to back up partition tables too.

Networking Getting connected

As mentioned, it is possible to set *Partition Image* up to work over a network. Then you can have a central server with your image files, making it much easier to centrally control the data that you back up and restore.

The *partimaged* program runs the server, and all you have to do is open a shell and run it. The user 'partimag' should exist on your server, so you have to create this before you run the server for the first time.

A screen will open showing the status of the 10 allowed connections, and that's it – the server is ready for operation.

Remember to make a note of the IP number if the server's hostname isn't included on DNS.

Bear in mind that the limiting factor of backing up over the network is liable to be your network speed. Backing up a 10GB partition over a 10MB line may take you more than a few hours.

The partition table is the information stored on the drive which tells the OS where, what type and how many partitions are on the drive. As you'll probably know if you've ever partitioned a disk, there are all sorts of ways to divide one up. Without the information on how a disk is partitioned, it's going to be a lot harder to restore images if you need to.

The *Partition Image* software itself does not (yet) save the partition information, so you'll have to save this yourself.

For a disk which has a few primary partitions, this is not too difficult. The original DOS specification for partitions allowed for up to four partitions on a single device. Under Linux these would be, for example, hda1, hda2, hda3, hda4. These types of partitions are known as primary partitions, and every disk will have at least one. The information describing these partitions is included in the MBR (Master Boot Record) of the device, which is always contained in the first 512 bytes of the disk.

But having only four partitions became a bit of a limitation, especially when much larger disks came along. The solution was to create logical partitions inside one of the primary partitions. This works by allocating a primary partition, then chaining the data for the logical partitions inside them. The first logical partition points to the location of the next and so on. That's why you may often find a disk with partitions labelled hda1, hda5, hda6, hda7, but with no hda2. The numbers 1-4 are reserved for the primary partitions.

Having these linked partitions is a bit of a pain because there is no single part of the disk to read the data from. So in order to save the partition info for the disk, we need to do two things:

- 1 Back up the Master Boot Record.
- 2 Back up the information for any logical partitions.

There are one or two utilities to help with this. If you are using Mandrake 8.x, you can use the *diskdrake* utility to store the partition information to a file for you. Select the drive and click on the "More" button in the bottom of the Window. This stores the data in a fairly human readable form too.

The alternative is to back up the data using standard Linux tools. The first task is easily accomplished using *dd*. As we are only talking about a 512-byte file, it's more than up to the task. For example, to save the MBR from your first IDE hard drive:

```
dd if=/dev/hda of=hda.mbr bs=512 count=1
```

To back up logical partitions, you can use *sfdisk*, one of the *util-linux* tools packages which will be installed as part of almost any distro. If you don't have it, you can get *util-linux* from your local kernel mirror or the main www.kernel.org site.

You should be careful with *sfdisk*, as it is possible to lose data using it. Check out the man pages for the options if you are unsure. To back up the partition information, use:

```
sfdisk -d /dev/hda > hda.sfd
```

Compression methods comparison

Find the most suitable compression for your data

Type	Time taken	Size
No compression	142s	81MB
<i>lzo</i> compression	144s	43MB
<i>gzip</i> compression	146s	41MB
<i>bzip2</i> compression	155s	38MB

These simple tests serve as a guide only. They were performed on a PIII/850MHz machine with standard IDE hard drives. The differences in time taken and final size will differ depending on the type of data stored in the partition and the speed of your processor.

Hard drives

Check your drive configuration for optimum speed

Before performing a disk intensive task like backing up or restoring a partition, it might be a wise idea to make sure that your drives are configured most efficiently. Normally, if you just install a Linux distro and don't check the settings, things will be set to the lowest common denominator. This works, but doesn't get the most out of your hardware.

The first and most basic thing to check is that your system is using an appropriate bus speed. You'll probably find that by default your IDE bus speed is set to 33MHz, which is a safe value for

PCI/IDE systems, but you may want to change this. The bus speed setting can dramatically effect the timings for pio modes if enabled on your drives.

You should also check whether your drives are using PIO modes, DMA or whatever. The *hdparm* utility can be used for setting these modes (it's part of the *util-linux* package you can find on the kernel website and its mirrors). However, you can really screw up your disk using this, so read the documentation! In doubt, best stick to slower backups...



This will dump the partition info into the file *hda.sfd*. Armed with these two files, you should be able to reconstruct the partition table of any disk. So for goodness' sake, copy these files somewhere safe. Please don't leave them on your hard disk – they won't be much use to you there if that's the disk or partition that has a problem! Best to keep them on a floppy (or a couple of floppies for extra security), or better still, burn them to the bootable CD you're going to make later...

The website contains up to date information, usage guides and many different versions of the software itself.

Backing up a partition

Before we start backing up a partition, you should first know where the partition is. *Partition Image* uses devfs-style partition labels, so if you are backing up *hda1* for example, it will be listed as *ide/host0/bus0/target0/lun0/part1*. For *hdb1*, the listing would be *ide/host0/bus0/target1/lun0/part1*. There is a little help in that *Partition Image* will also display the filesystem installed and the partition size. Remember that the partition should really be unmounted before backup, to ensure that no files are open or can change during the backup.

Restoring partition information

If your partition information has been lost, or you want to duplicate the partitions from an identical drive, then we can restore the information we saved earlier:

```
dd if=hda.mbr of=/dev/hda bs=512 count=1
```

```
sfdisk -O changelog.sfd /dev/hda <hda.sfd
```

IMPORTANT: This will rewrite your partition table. Make sure that you type the commands correctly and use the correct backup files. Incorrect partition data could result in loss of data.



TutorialPartitionImaging

« The *sfdisk* utility has safeguards to make it hard to write partition data to a mounted drive. You should unmount the drive before you write the partition table. Unfortunate things can happen if you change the partition info when files are being accessed on the disk! If you use the wrong file, or discover a mistake, you can usually recover the partition with the aid of the changelog file we created when we were writing the *sfdisk* data:

```
sfdisk -l changelog.sfd
```

Of course, if you have used *Diskdrake* to back up your partition information, you can also use that to re-write the data.

Whichever method you use, you will normally need a reboot for the changes to be noticed by Linux, so do that before attempting to restore any partitions.

Restoring the partition

Once your partition tables are sound, you can begin restoring the partition. This is an even easier task than backing up because

there are less options to worry about. Select the partition you want to restore from the list which is shown (choose carefully) and enter the filename of the file containing the data. If you are restoring over a network, you'll have to select the IP or hostname of the server too.

Press F5 and, if you entered a description for the saved file, this text will now be displayed. You can 'simulate' the write, which at least will check for potential problems with the image file (if it has become corrupt, for example, or if it's the wrong size).


You can also choose to erase empty blocks with zero values. This will overwrite any data in the partition and make it all but impossible to retrieve any of the data from the drive.

Conclusion

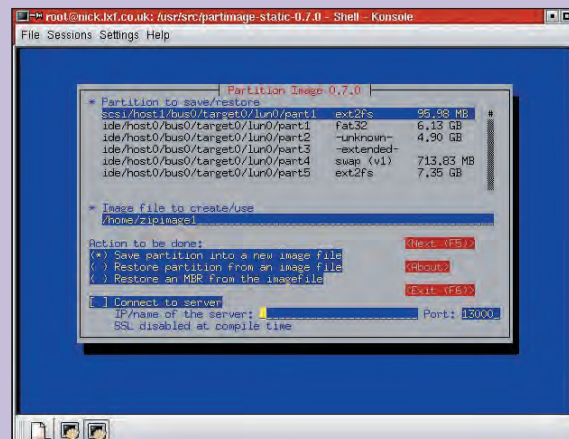
While not exactly fun, *Partition Image* certainly makes the backing up of partitions fairly straightforward and simple. Nobody likes backing up data, but backing up complete partitions offers an easy and convenient solution for those who find it necessary. And after all, if you have all your partitions backed up, what's the worst that can happen? ■

Partition Image step-by-step

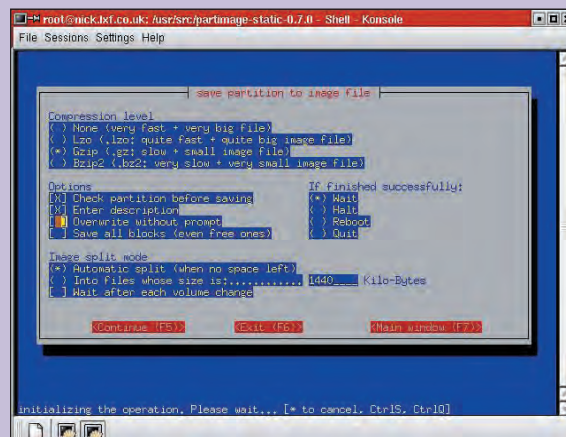
We walk you through the backup procedure




STEP 1 Run the *Partition Image* program from the console. Make it a reasonably sized window so you can see clearly what's going on! You'll now see a curses-style display of the main menu. The first detected partition will be highlighted. Use the cursor arrows to highlight the partition you wish to back up.



STEP 2 Press the TAB key to move between the options. Type in the pathname of the file to store the partition data in. You can leave the rest of the options for the moment, as we are not connecting to a server, nor are we performing any restore options (the default is set to Backup Partitions, as indicated by the asterisk).



STEP 3 Press F5 to continue: On the next screen we can set the backup options. The first bank of options selects the compression level. The default options are usually fine. If you choose to check the partition (a good idea) you'll soon find out whether you remembered to unmount it first.



STEP 4 If you want to split the files to fit on CDs or removable disks, you should check these options here. The Wait after each volume change can be useful if you are writing directly to removable media. In the example here we have specified a size suitable for fitting on Zip disks.

Creating a boot CD for x86

Get to grips with El Torito

Available from the *Partition Image* website is a useful El Torito image for creating a bootable CD. There are a number of potential uses for this:

- Backing up your computer's root partition (which you won't be able to do when it's mounted)
- Backing up a partition from a non-Linux machine (ie a Windows desktop)
- Restoring partitions to a machine with no viable OS


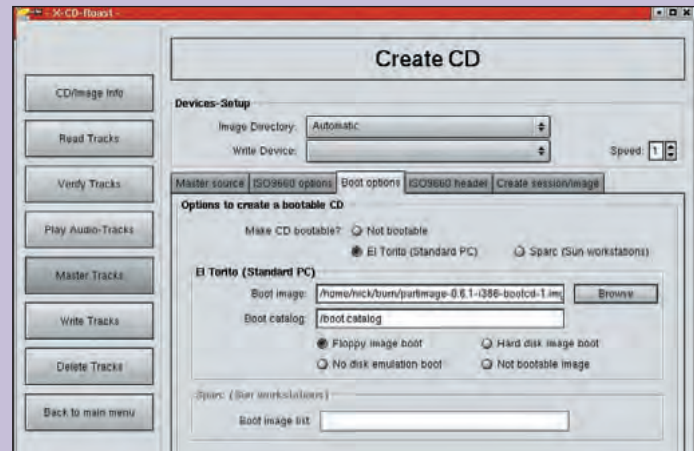
Creation of the boot disk is simple. You just have to copy the El Torito image to a directory and master a CD, using this file as the El Torito image. (For various reasons, this is the name given to the file

which is used as the bootable portion of a CDROM.)


If you have a generic partition image, you can burn this to the CD too, enabling you to boot and restore your system without access to any other drives. As the El Torito image takes up a negligible amount of space on the CD, there is 650MB, more or less, of space available for any partition backups.

You can create an image and use it with the standard Linux tools, such as *mkisofs* and *cdrrecord*, but as GUIs like *XCDRoast* also support El Torito information now, it's probably easier to use this instead.


Point *XCDroast* at the El Torito image to create a bootable CD.




STEP 5 Hit F5 and you will have the option to enter a useful description. Do so and then the operation will initialise. This is the point at which the filesystem is checked (if you chose that option) so it can take a few moments. All you have to do is wait patiently.



STEP 6 An info screen will pop up to tell you about the partition you are backing up. This may have various pieces of data, depending on the filesystem in use on that partition, in addition to some info about blocks and the size and used space on the partition. Press ENTER to continue.



STEP 7 The operation will now start, giving you a running total of the amount of data stored so far, and various other statistics to keep you happy. The time taken for saving the image obviously depends a great deal on how much data you are backing up!



STEP 8 When the operation has finished, you'll see this screen showing you what exactly has happened (note that if you have chosen to wait for split files, you'll get messages in between for media changing). Click on OK and that's it! Your data is backed up.

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!

Our experts

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.



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



CDROM from Linux

Q I've been picking up your fine mag on and off for about a year. I've always found an interesting article or two and something useful on the CD. I especially like Mandrake Linux 9.0 on the Christmas 2002 issue. Just one question, though.

I have a Dell Latitude LT laptop. The CDROM is connected via a KME KXLC005 PCMCIA SCSI controller. What driver do I need to support this? The installer-supplied drivers all result in the message "No CDROMs found". I worked around this by copying the CDs to my hard disk and installing from there, but I REALLY want to use the CDROM from Linux (and be able to get rid of Windows completely). Many thanks for your help!

A The SCSI chipset you are using would appear to be supported by the `qlogic_cs.0` module, although the PCMCIA documentation does not specifically list the controller you are talking about. You've not supplied any system logs or debugging information – having those available would make it much easier for us to diagnose the problems you are looking at.

Since Mandrake 9.0 is a somewhat up-to-date distro, and the chipset you are using is reasonably old, it's probably fair to assume that it's likely that it isn't supported. Depending upon your laptop, you may wish to switch to a USB CDROM drive, or find a cheap PCMCIA SCSI adaptor to replace the one which does not work.

Information on the cards currently supported by PCMCIA is available at <http://pcmcia-cs.sourceforge.net/> **ftp/SUPPORTED.CARDS**, if you wish to check out the supported chipsets before making a purchase.

terminal window: david@macha:~ (pts/18)

```
maeve:~# repquota -a /home
xxx Report for user quotas on device /dev/maeve/home
Block grace time: 7days; Inode grace time: 7days
Block limits      File limits
User      used    soft    hard    grace    used    soft    hard    grace
-----
root      --    13         0         0         4         0         0
debian    --    29         0         0         9         0         0
david     -- 5881991     0         0    212568     0         0
chowells  --   7436     0         0      888         0         0
mccgoldrick -- 139         0         0      23         0         0
mccgoldrick -- 1639         0         0      68         0         0
melanie   --    26         0         0      11         0         0
msutton   -- 3524         0         0     92         0         0
cbutcher  --    30         0         0     13         0         0
dbrooke   --    37         0         0     10         0         0
jbruce    --    21         0         0      7         0         0
pcraig    -- 7491         0         0    1134         0         0
deana     -- 4333         0         0     152         0         0
chris     --   644         0         0     14         0         0
nheap     -- 3724875     0         0    11997         0         0
maeve:~#
```

Since users often don't listen to what admins tell them, quota support avoids having a user consume large amounts of storage space.

Perl text solution

Q Please help me with this nightmare. I have text files from a program that I need all the numbers updated. The files has seven or eight hundred records each that begin with something like 0266666660004100 01/15/0309:49. The first part is the part I need to modify. The first two digits are the site code. The next six digits (the 6s in this example) are a number that I need to change for the first line one to a number I select and have the rest of them autoincrement to the end of the file. The third part is the account number that I have to keep. There is plenty of data to the right of the date field. I figured it would be something like `perl -e (lots of magic here) filename 777777`. It would make the first record's number in the file change from 666666 to 777777 and then autoincrement the next record which is on a new line to 777778, to the end of the file. I know this is possible because it is Perl but I have not a clue as how to begin. Please help.

Fred Howell, via email

A The quickest way to do all of this is with a regular expression and a nice little loop with Perl.

```
#!/usr/bin/perl
$n=$ARGV[0];
open FILE, "stuff.txt";
open OUT, ">out.txt";
while(<FILE>) {
    /^(0-9){2})(0-9){6})
    (.+)$/;
    print OUT "$1.$n.$3\n";
    $n++;
}
close FILE;
close OUT;
```

This doesn't do any checking of what you feed it, so it'll happily dump anything into it which you pass via the command line. It's probably quite wise to test this with a backup of your extra-special data, since we only tested it with half-a-dozen fields, but a loop of this type should work with a massive number of lines.

Opera blues

Q I am running kernel 2.4.19 – Pentium III – Opera 6 – Linux-only Desktop with no firewall or special security

precautions beyond the usual passwords etc. Recently *Opera 6* crashed and the machine hung whilst on the Internet. Subsequently I found a **ptrace()** report in the system log, never having seen one before in 18 months of running Linux. After the initial problem *Opera* would not bring up a screen but spent its time creating spurious directory files which presumably filled the available ext2 inodes and again caused the system to hang. Clearing out *Opera* and all related files seems to have cured the problem. Having read of past security breaches related to **ptrace()** I wonder if the latest kernel is again vulnerable. I should add that 'crashes' have been notable by their absence until this event. Occasionally KDE or the odd program would give trouble but it has been possible to 'kill' the particular application and continue without the need to reboot.

As this is a security question I would be happier if you did not use

my surname should the question appear in *LXF* – no point in issuing a challenge or inviting trouble!

Peter, via email

A To our knowledge, the **ptrace** exploit has been fixed in 2.2.15 and above kernels. You didn't tell us what the specific **ptrace** call was in the kernel logs, so we can't provide any hard and fast information about the problems you have been experiencing. It's worth noting that the earlier **ptrace** issues were for local root exploit only, so someone would have needed to crack your box and obtain a user account on it, or at least be able to execute binaries on your system, before being able to get root. With a wide open box without a firewall, depending what network services you are running, it may or may not be easy to obtain a shell on your box. You said you've been running Linux for 18 months, so have you been keeping up with security patches since then? Have you updated your distribution with the latest security releases and all that fun stuff?

You really want to run a root kit auditing program, such as **chkrootkit**, on your box to see if someone has exploited it or otherwise obtained root on it through a rootkit. You really need to setup a firewall if you've got a box sitting on the Internet, and since it's only about half a dozen lines with **iptables**, if that, it's really easy to do:

```
iptables -A INPUT -i ppp0 -m state
--state RELATED,ESTABLISHED -j
ACCEPT
```

```
iptables -A INPUT -i ppp0 -m state
--state NEW -j LOG
```

```
iptables -A INPUT -i ppp0 -j DROP
```

Of course, **ppp0** should be replaced with the appropriate external interface on your system. Based on the fact that your box is possibly exploited, and if you don't have time to do a complete security audit, I would recommend copying the data you need somewhere safe, such as a CDROM, then completely obliterating the current install and reinstall the most up-to-date version of your favoured distribution that is available, then install the latest patches and updates for it.

Muted laptop

Q I have installed RedHat 8.0 onto a Toshiba Tecra 8000.

It all seems to work well, except for the sound. I followed the instructions at the bottom of this page: www.shroom.com/linux/laptop/8000-2.html, but my box has:

```
WSS I/O=530H
SBPro I/O=220H
Synth I/O=388H
WSS/SBPro/MPU402 IRQ
level=IRQ5
WSS(Play) DMA=1
WSS(Rec) DMA=0
Control I/O Address=538H
MPU401 (MIDI I/F)=330H
```

The example given has Control Address=0x370, not 0x538. All three of my Toshiba Tecras have the same settings, with no obvious way to change the control address to 370H.

I can enable sound, as root, system-wide, by using one of several different soundcard choices in **sndconfig**, but sound is lost again >>

A QUICK REFERENCE TO: GPHOTO2

Digital cameras are particularly popular, now that higher-end models are falling in price, and a entry-level camera can be bought for a few hundred pounds, if that. Of course, a digital camera isn't a whole lot of use unless we can copy the images from it to our workstation or laptop.

Nearly all digital cameras have a USB connection, so it's reasonably straight forward to setup the hardware side of it all. Using the standard USB modules for Linux, once we plug our camera in and power it up, something exciting should happen and **dmesg** should note that our camera now exists.

```
hub.c: new USB device 00:1f:2-1,
assigned address 5
usb.c: USB device 5 (vend/prod
0x4a9/0x3048) is not claimed by
any active driver.
```

Fortunately, to get the data off our camera we don't need any kernel modules or patches, as everything is done in user-space via the **/proc/bus/usb** interface to

the USB subsystem. All we need to do is use the **gphoto2** package to pull the contents of our digital camera to our hard drive.

However, before we can do this, we need to chown the **/proc/bus/usb** device so that a non-root user can access it. Since it said 'device 5', then we need to chown **/proc/bus/usb/005** to someone useful. This can be achieved automatically using **hotplug** with the scripts which you can find within the **gphoto2** documentation.

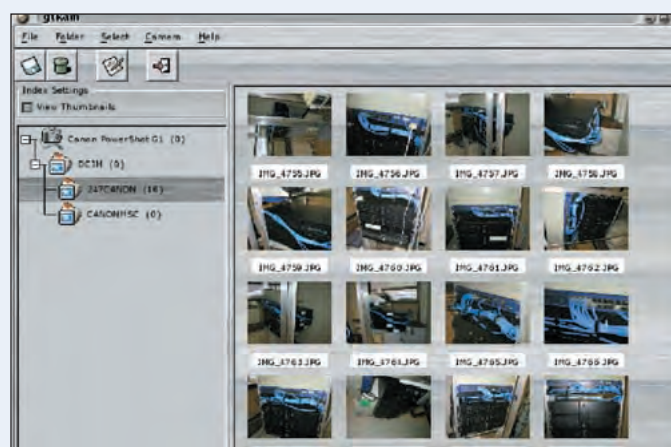
Now we've got our camera connected, we need to tell **gphoto2** what we want the computer to do with it. **gphoto2** will be able to automatically figure out the model of our camera, so we can simply ask it to pull all the images:

```
gphoto2 -- auto-detect -P -R
```

which will dump the contents of the camera into the current directory. We can then delete the files from the camera with:

```
gphoto2 -- auto-detect -D -R
```

Both of these will recursively go through each folder on the camera



gtkam is a nice little front end to **gphoto2**, which allows us to quickly grab images from our camera.

and download the contents. We can get a list of folders with:

```
gphoto2 -- auto-detect -l
```

Once we've decided which folder we want to download, we can pull that one down from our camera with:

```
gphoto2 -- auto-detect -P -f
<folder>
```

Of course, if a console based program is not to your tastes, one can use **gtkam**, which is a front-end to **gphoto2**. This is a very simple application which

downloads the thumbnail images from the camera, which then allows us to select specific images and pull them down our cable to the system. There is also a nice plugin for **The GIMP**, which will allow us to browse photos then pull them down directly into **GIMP** without having to go via the hard drive. There's lots of documentation on **gphoto2** at www.gphoto.org/gphoto2/, and it is available as a package for most distributions.

FREQUENTLY ASKED QUESTIONS SYSTEM ADMINISTRATION

FAQ How do I create a new user with Linux?

The standard way to create a user with Linux is with the **useradd** command, although most distributions have their own variant, such as **adduser**. **useradd** has a very simple syntax:

```
useradd [-u uid [-o]] [-g group] [-G group,...]
        [-d home] [-s shell] [-c comment] [-m [-k template]]
        [-f inactive] [-e expire] [-p passwd] name
```

So, if we wanted to create a user called 'david' who's default group was 'users', we just need to do

```
useradd -g users david
```

The group details are a little complex, as a user has both a default group, along with a list of other groups which they can switch to with **chgrp**. One thing to note with **useradd** is that the **-p** option inserts the supplied text directly into the `/etc/passwd` file. If we want to actually shove a working password straight in, we need to create a password hash using **crypt()**, with something like Perl or C.

FAQ Can I give them a default password?

If we're not looking at using the **-p**

flag, it's easy enough to use the **passwd** program to set a password for a user. For the user 'david', we can simply do this:

```
passwd david
```

which will ask for a new password. Note that this should be performed as root, otherwise it will complain.

FAQ Can I distribute authentication information over a network, so I don't need to make every user on each system?

Well, if we've got lots of free time or are being paid by the hour, we can. Unfortunately, life is not always that easy, so we can use a system known as NIS to distribute authentication information. NIS is a server/client system, so we run a NIS server which contains all of the authentication information, then on each system we run the client which permits the system to do authentication lookups via the server.

Information on NIS can be found at <http://en.tldp.org/HOWTO/NIS-HOWTO/>

FAQ Can I stop users logging into a particular

system, but still have their user information available?

When using NIS, we can modify any `/etc/passwd` information, such as the shell. To only permit users from the 'admin' netgroup to shell into the host, we can do:

```
+@admin::0:0::
+*:0:0::/bin/false
```

FAQ How do I set storage quotas for each user on /home?

See the tabular lines of code in the box below. Most file systems, including ext3 and reiserfs have quota support, which requires us to compile quota support into the main kernel. By mounting the `/home` file system with the 'usrquota' option, we can give each individual user their own storage quota, which is based on both bytes used as well as inodes consumed.

```
maeve:~# repquota -a
```

```
*** Report for user quotas on device /dev/maeve/home
```

```
Block grace time: 7days; Inode grace time: 7days
```

User	Block limits			File limits				
	used	soft	hard	grace	used	soft	hard	grace
root	—	13	0	0	—	4	0	0
debian	—	29	0	0	—	9	0	0
david	—	5881991	0	0	—	212568	0	0

The command **repquota** displays the current quota information for all users on the system:

As we can see, there is both a **soft** quota and a **hard** quota. A **soft** quota can be exceeded, but the user is emailed every day about being over the limit. A **hard** quota is fixed, and the user is unable to write any more data than this limit.

For each user, we can run 'edquota <username>' and modify the file as appropriate. If lots of users are going to have the same quota information, we can start by creating a prototype user quota, then replicate it with this:

```
edquota -p <protouser> david
```

FAQ Can I limit the number of processes or memory my users can consume?

ulimit can be used to limit the number of processes a user can have running, and the total amount of RAM

after a reboot. The new Red Hat sound configure gives me 'no soundcard present' (or something like that). What, must I do/edit, to get sound working, on a permanent basis? Thanks for any help.

Gormb, from the forum

As the system uses the OPL3-SA2 driver, you should be able to use a configuration option in `/etc/modules.conf` similar to the following:

```
alias sound-slot-0 opl3sa2
options opl3sa2 io=0x538
mss io=0x530 mpu io=0x330
irq=7 dma=0 dma2=3
options opl3 io=0x388
```

Since you're at least able to get it to do something with **sndconfig**, you may want to run **dmesg** afterwards

and figure out exactly what it decided to load for you. However, the **op3sa2** module should be able to autoprobe for your sound card, so simply doing **modprobe opl3sa2** should allow it to find your sound card automatically. **sndconfig** should save your options following a reboot, so you should check to find out if it's writing anything useful to `/etc/modules.conf`, much like the example giving above.

Not on display

I have tried several times to load Mandrake Linux 9.0 and it always fails at the same point, which is the selection of the display. Whether I select the correct monitor (Mitsubishi Diamond Scan 50) or a generic 800x600, I get a screen with

vertical colour bars with an "Is this correct Y/N" dialogue box and whichever option I take – or leave it to time out – the screen goes to more spaced and "dotty" colour bars and the machine locks.

I am running an Athlon 1600 XP with a K755A motherboard. IDE-0 and 8GB of IDE-1 are Windows98 but Linux is loaded into the default partitions of the rest of IDE-1, which is an 80GB Deskstar. The video card is PCI because I had problems with an older AGP card when I installed the new motherboard. At one stage I tried to run rescue but that did not get anywhere so did a new install and formatted the two partitions that the installation program selected.

If anyone has any suggestions, I

would be most grateful. The immediate task is to check out an IBM RPG emulator that I found on the Net.

Peter Kinsman

It sounds as if Mandrake is picking the wrong video card, so the monitor isn't getting the correct signal. You neglected to mention what video card is installed or the chipset it uses, as well as the driver Mandrake decided your video card needs to be using. Since the box locks up, it's not going to be an issue with the monitor getting the wrong signal, since you can send the monitor anything you really like, and even if it doesn't like it, your video card isn't going to work.

Once you find out what video card you've got, go to www.xfree86.org

to be used by those processes. An example usage of **ulimit** would be:

```
ulimit -u 35
ulimit -s 2048
ulimit -c 0
ulimit -m 25600
```

The first option limits the number of processes to 35. It's best not to set this too low, otherwise the user won't be able to actually do anything. **-s** limits the stack size to 2Mb and **-c** sets the maximum core size, which is the dump following a segfault, to zero, meaning that no core file will be generated. Finally, we limit the amount of RAM used by the user to 25600Kb.

FAQ What about CPU usage? Can I limit the amount of CPU time a user or process can consume?

'nice' can be used to limit process priority at runtime, and 'renice' can be used to modify a running process priority. However, this is all open to abuse by users. Instead, we can apply the 'fair scheduler' patch which reduces process priority when a particular user is using too much CPU, effectively creating a dynamic 'renice' system for the processes. This patch can be found at <http://www.surriel.com/patches/>

and figure out what driver it needs to be using, and you might be able to convince Mandrake to use it.

Does it probe for the video card correctly and pick an driver, or did you have to manually select one? If you have further problems, it may be worth posting to our forums all the information on your video card and X installation so that more useful information can be made available to you by those with similar experiences.

Missing modules?

While setting up my Zip 100 USB to work with MDK 9.0, I came across the tutorial here: <http://linux-sxs.org/usbzip.html>. But step 7 of the tutorial says "If you want these modules to be available the next time you boot add these lines to

/etc/modules/default:

```
sg
ide-scsi
et cetera.."
```

My problem is that in my Linux installation there is no modules subdirectory of etc so I can't find the file 'default'.

Can you help me please?

Tung Tse, London

The instructions you were following were for a distribution other than Mandrake, so /etc/modules/default is the correct location on that distribution, but not with Mandrake Linux. You should put all of that lovely stuff in /etc/modules, which is the appropriate place with Mandrake. The rest of the documentation is applicable to all recent Linux distributions which have the required modules available.

Proxying guidelines

I need some help finding (tutorials and architecture guidelines) for configuring Linux-based proxies for more advanced uses like: Reverse Proxying, Proxying SSL Authentication in conjunction with PKI. Any suggested products or evaluations would be greatly appreciated too.

Joe, via the forum

Apache using mod_proxy and/or mod_rewrite allows you to do reverse proxying of either HTTP or HTTP-SSL connections. You may wish to read the documentation available at http://httpd.apache.org/docs/mod/mod_proxy.html, http://httpd.apache.org/docs/mod/mod_rewrite.html and <http://httpd.apache.org/docs/misc/rewriteguide.html>. On the PKI front, you may wish to check at www.modssl.org to look at the authentication options it adds to Apache.

There are a number of solutions which supply these capabilities, such as those from www.rainbow.com, which supports all of the options you desire. Of course, depending on your knowledge and your budget, these may not be appropriate.

Bloat-free WM

I use Fluxbox all day every day on Red Hat 8.0, and it works really groovily, but I'd like a desktop similar to KDE or

gnome, ie easily addable shortcuts etc. I've tried and not really got on with iDesk. Is it possible to use just the desktop of KDE or GNOME with Fluxbox without the complete bloat of either of them?

Failing this, do you know of an equivalent? I've tried the usual Google, Freshmeat, SourceForge etc. I know my demands are a bit moon-on-a-stick...

Crispin Veall, Somerset

You can run the KDE 'kfm' toolbar, or GNOME's 'gmc' within any Window Manager you like. All you need to do is run the corresponding program from your .xsession file, so that it starts up when you login to the system via X. Both of these may have problems, depending how your chosen window manager handles the background, but on the most part, they should work just fine.

Move disc space

I am running a very old version of Red Hat (5.1) It works fine as an MRTG, Nagios, Apache, CHRONY server (plus a few more) unfortunately, I am coming to the point where I need to change the distribution of disc space – currently my file systems look like this:

```
[root@homer etc]# df
Filesystem 512-blocks Used Available Capacity Mounted on
/dev/hda1 8104264 69920 7615052 1% /
/dev/hdb1 93264 1172 87276 1% /boot
/dev/hdb7 3574200 30 3389426 0% /opt
/dev/hdb6 4071150 836630 3024076 22% /usr
[root@homer etc]#
```

As you can see, /usr is filling up, but /opt is empty. Is there anyway for me to move space from /opt to /usr?

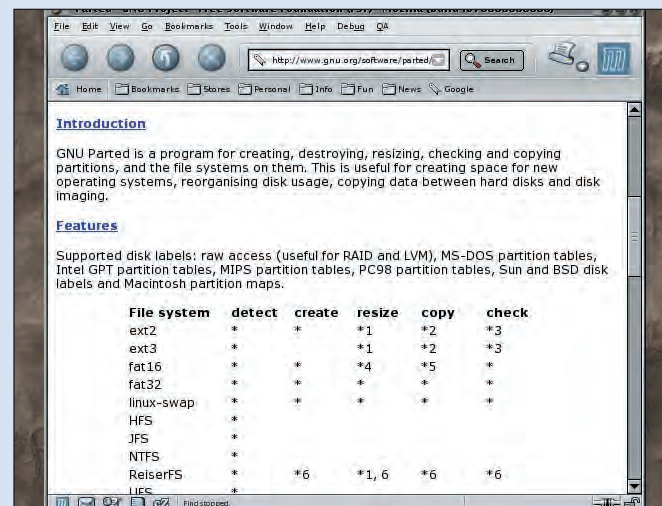
Please remember, this is a 'headless' pc, so I can only use command line! Many Thanks!

Paul Vallack, via email

Command line? No problem! The easiest way to do this is to use a command known as parted to remove the hdb7 filesystem and then resize hdb6. I'm assuming that hdb6 is found first on the disk, which makes life much easier, since we can just resize it, rather than having to move it around on the drive.

parted can be found ready for download at <http://www.gnu.org/software/parted/parted.html>, or on the cover disc of the latest issue of LXF.

Note: to perform this operation both /opt and /usr will need to be unmounted, so you may wish to be looking at setting up a serial console on the system, so that you can access it without having telnetd or sshd binaries available.



parted, while not being at the level of PartitionMagic, allows basic partition functions to be performed.

Answers



Remote mail

Q I have several employees who travel around the country and need to be able to send mail via SMTP from anywhere via the Internet. However, sendmail is configured so that users cannot communicate with any IP addresses not our LAN. How do I let these remote users send mail without experiencing large security risks by opening my mailserver to viruses or spammers? The server runs RedHat 7.3 and we use m4 to generate sendmail.cf

A What would probably suit your circumstances best is SMTP authentication. This will allow users to send mail through your server no matter where they originate from provided they authenticate first. How do you do this? Make sure that the following lines appear in sendmail.mc:

```
define(`confAUTH_OPTIONS',
`A')dnl
TRUST_AUTH_MECH(`DIGEST-MD5
LOGIN PLAIN')dnl
define(`confAUTH_MECHANISMS',
`DIGEST-MD5 LOGIN PLAIN')dnl
```

Backup your sendmail.cf and generate a new sendmail.cf with m4:
m4 /etc/mail/sendmail.mc > /etc/sendmail.cf

Make sure that /etc/pam.d/smtp exists and contains:

```
auth required
/lib/security/pam_stack.so
service=system-auth
account required
/lib/security/pam_stack.so
service=system-auth
```

Next, you should verify that /usr/lib/sasl/Sendmail.conf exists and contains the following:

```
pwcheck_method:pam
```

Restart sendmail and that should complete the configuration. All that remains is to add the user accounts for each employee who will be working remotely by typing **saslpasswd <username>** and entering the password for that

user. Don't forget to configure the mail client so that it knows to authenticate too.

Extract from RPM

Q I would like to extract files directly from an RPM without having to install the program. I've read the man pages for RPM but I can't find any reference to this. Is it possible to do this? Thanks in advance.

A Yes this is actually possible if you run the following command:

```
rpm2cpio packagename.rpm | cpio
-id
```

This will create the entire directory tree for the RPM, so it's best to run this from a temp directory.

Hidden pages

Q Here is my problem. I want my website to be accessible to the public. But still allow

certain pages to be accessible only to authorised users. How do I make Apache ask for a username and password? Also, do these users have to be users on my system?

A Apache can do exactly what you are looking for! By using .htaccess files the users do not need to have accounts on your system. You can also protect files on a per-directory basis, but to do this you'll need to store all the files you want to protect in this directory

First, make sure that your httpd.conf (usually in /etc/httpd/conf or /etc/httpd) contains the following directive for the directory which you want to protect:

```
<Directory /var/www/httpd/
privatefiles>
AllowOverride All
</Directory>
```

If you are doing this for a virtual host then make sure that this appears within the **<VirtualHost>** and **</VirtualHost>** tags. Next you need to create the .htaccess file in the directory you want to password protect. This file contains the authentication information and also points to the file which contains the encrypted login details, it should look as follows:

```
AuthUserFile /var/www/.mywebsite.
htpasswd
```

AuthType Basic

AuthName "Members"

require valid-user

The **AuthUserFile** can point to any place on the filesystem which Apache's user can read and it is recommended that you do not store it in a directory that can be accessed from the web, that way there is no possibility of somebody guessing the filename, downloading it and attempting to decode it. In the above example I've assumed that your document root is /var/www/html. Now you can add users to the .htpasswd file by running the following commands:

For the first user:

```
htpasswd -c /var/www/.mywebsite.
htpasswd username1
```

For each further user you add:

```
htpasswd /var/www/.mywebsite.
htpasswd username2
```

The system will prompt you to enter a password for each user. The only thing left is to restart Apache and test your configuration.

Spam, spam, sp...

Q I run a small mailserver for some of my friends and myself. The server runs sendmail 8.11.6. Recently I've been getting more and more spam, in fact it's reached the point where I probably get more spam than real

About Rackspace Managed Hosting

Rackspace Managed Hosting™ provides complete managed hosting solutions with Fanatical Support to small to medium-sized customers as well as large enterprises. All customer platforms include state-of-the-art data centres,

customised servers, burstable connectivity, 99.999% uptime SLA, a dedicated account manager, instant emergency response and access to live expert technicians 24x7 for support of all hardware and core software.

Founded in 1998 with locations in San Antonio, Texas and London, the company manages over 7,000 servers for customers in more than 80 countries. For more information, visit www.rackspace.co.uk

HANS HUBERLAND RACKSPACE MANAGED HOSTING EUROPE

www.rackspace.co.uk

Hans Huberland is the resident

Linux expert at Rackspace Managed Hosting, the only profitable managed hosting provider in the UK.

Hans joined Rackspace Managed Hosting in early 2001. A Level 3 System Administrator, Hans and his team currently manage and support over 7000 servers, each with their own configuration. As an ardent Linux advocate, not only does he work around the clock to ensure clients are stable and that problems are resolved quickly, he also spends much of his spare time learning about Linux. Since first becoming interested in

Linux while at college, Hans has dedicated his time and education to becoming an expert in the Linux arena and in particular, the system administration of web and mail servers. He honed his skills while working for computer component distributor, Axiz and small startup ISP/hosting company, Tinroof, before joining Rackspace two years ago.

Over the coming months Hans will be answering questions relating to system administration using Linux. Email your questions to sysadminqa@rackspace.co.uk and you could win a Linux PDA!



mail. I've been told by a knowledgeable friend to set up a spamhaus.org blacklist. How do I set this up and stop this invasion?

A You'll need to make the following changes to your sendmail.mc file and then write the changes to sendmail.cf using m4.

```
FEATURE(`dnsbl', `sbl.spamhaus.org',
`550 Email rejected due to
spamhaus.org blacklisting')dnl
```

The first field builds DNS based blacklisting into the sendmail.cf file. The second field holds the server you want to use for lookups, and the last field is a custom message to send back to the spammer. Only the first field is required, sendmail will then default to using www.mail-abuse.org which is a subscription based service but they do offer free accounts to hobby sites. Check out their site for more information.

Samba address

Q I have a machine that I am trying to run Samba on but something else is hogging port 119. On my machine, the /var/log/smb.log says "address already in use". I have already checked in /etc/services and tried netstat -a. But I still can't track the program down. Please can you offer me some help?

A Try running netstat -p, this should give you the PID and program name of the program, which is bound to that port.

Form file uploads

Q I seem to be having some trouble with forms and PHP. I keep having problems with the following HTML in PHP:

```
<form action="/path/to/script.php"
method="post"
enctype="multipart/form-data">
```

The scripts work correctly when using method="post" and method="get" but not when the enctype is added in. Any thoughts? The above works on other servers.

A This is most likely caused by "file uploads" turned off in your php.ini file. When this setting is off, php discards any form with type multipart/form-data. You'll need to set this setting in php.ini and restart Apache.

Creating space

Q My company's ageing file server has recently become far more utilised than expected. This has led to the server crashing every couple of days. I've been monitoring utilisation and we appear to be running out of swap space. I've already partitioned all the available disk space so I can't create another swap partition.

Is there any other way to decrease the size of one of the other partitions and use the newly created space for more swap? The server is being replaced in the next few months but I'd like it to be stable until then if possible. The server's disks are in a RAID array.

A Swap space does not have to be a partition. Rather than trying to change partition sizes you can allocate unused space on one of your partitions to swap. First create a file of the size that you wish to assign for swap (512Mb in this example):

```
dd if=/dev/zero of=/swapfile
bs=1024 count=524288
```

This command creates a file called /swapfile containing 524288 blocks of 1Kb each (512Mb total) and fills it with zeroes (/dev/zero). You will need to initialise the swap area and then run sync to make sure the filesystem buffers are flushed:

```
mkswap -c /swapfile 524288 &&
sync
```

You'll need to change the permissions of the file to -rw- by running:

```
chmod 0600 /swapfile
```

Finally you'll need to enable the swapfile. To manually turn the swapfile on you'll need to type:

```
swapon /swapfile
```

Preferably you'll automate this from /etc/fstab. The settings should be the same as your current swapfile except for the devicename, where your current swap device might read /dev/sda2. You should make the new one /softfile.



missed one?

LINUX FORMAT BACK ISSUES

Every month *Linux Format* brings you the best tutorials, the essential reviews and the latest news. But if you've missed out on a must-read feature or a vital programme from our expertly compiled CDs and DVDs, order your back issue NOW! And remember, you need never miss an issue of your favourite Linux mag, if you subscribe to *Linux Format* (see overleaf for more details).



February 2003

Product code:
LXFB0037(cd)
LXFD0037(dvd)

DVD HIGHLIGHTS:
IPCop, Postnuke, Llypynd, Squid, Allegro, Cocoon, Fandango, gPhoto, Xrmap, Ted, AdvancedPHPdebugger, Flightgear updates

CDs HIGHLIGHTS:
Yoper, KOffice 1.2.1, Parted, DOM Tooltip, Guarddog, GL-117

MAGAZINE FEATURING:
Lindows on test, Serial ATA, Window Managers roundup, Arkeia review, Xft2, build your own Linux backup box, Java compiling, Jool Kwartz review



January 2003

Product code:
LXFB0036(cd)
LXFD0036(dvd)

DVD HIGHLIGHTS:
Evolution PIM/calendar/email client, Freevo personal video recorder, OpenZaurus, PilotLink

CDs HIGHLIGHTS:
Racer, Freedud, Linux From Scratch, GNOME Meeting, Opera 6.1, Film GIMP, phpOpenTracker, Tuxpaint, Snort

MAGAZINE FEATURING:
Optimisation tips and tricks, Netbox Cubit review, build your own desktop, Xandros, Cubit, CORBA, Realsoft 3D, Instant Messenger roundup



Christmas 2002

Product code:
LXFB0035(cd)
LXFD0035(dvd)

DVD HIGHLIGHTS:
Mandrake 9.0 - 3 CD download edition, jEdit, BastilleLinux, Phoenix, VideoLAN, KDE 3.0.4

CDs HIGHLIGHTS:
3CDs, including 2CD Mandrake 9.0, SpamX, Stellarium, Krename, ImagePress, Nogger, DigitalMusicCentre

MAGAZINE FEATURING:
Huge C++ IDE RoundUp, LXF Awards nominations, gadgets for Xmas, User Mode Linux, compiling code, process accounting



December 2002

Product code:
LXFB0034(cd)
LXFD0034(dvd)

DVD HIGHLIGHTS:
KDE 3.1 beta2, Scribus, Smoothwall2, Movix, GNUPrivacyGuard, GCompris, Drip

CDs HIGHLIGHTS:
Unreal Tournament 2K3 (Demo), DemoLinux, Cardfile, RUTE, OOOictionaries, Mozilla binary packages

MAGAZINE FEATURING:
Distro reviews of SuSE 8.1, Red Hat 8.0, Lycoris and a mini-distro RoundUp, scanning, PHP acceleration, Linux Expo UK 2002



November 2002

Product code:
LXFB0033(cd)
LXFD0033(dvd)

DVD HIGHLIGHTS:
Debian 3.0, mFighter, OpenOffice.org (bugfix release), DVDrip, Cyrus IMAP Server

CDs HIGHLIGHTS:
Kylux 3, GNUCash, BXPro, KDevelop, Opera, Vega Strike, Parted, AnjutaIDE, GTransferManager

MAGAZINE FEATURING:
Sun's move into the Linux server market - with LX50 review, The Liberty Alliance, Systems programming, using OpenOffice.org, Homebase



October 2002

Product code:
LXFB0032(cd)
LXFD0032(dvd)

DVD HIGHLIGHTS:
Knoppix, Drip, Squeak, extra FlightGear maps, Ogg Vorbis 1.0, Knoppix

CDs HIGHLIGHTS:
FlightGear (runs from disc), Aglaophone, UAE, Clam Antivirus, Perl 5.8, Quanta Plus, Netclipboard, Mah-Jong, HTML-Mason, WebSuck, Epsutil

MAGAZINE FEATURING:
Building better databases, 'Trusted Computing' - beware Palladium, USB 2.0, firewall roundup, Amiga emulation, Gentoo review

Other back issues also available, please call for further information, giving the issue number and/or date required

To order, call the back issues hotline on **0870 4448645**

or **+44 1458 271178** for overseas orders

Remember to quote the issue code shown above and have your credit card/switch card details ready.

Prices (inc post and packing): CD - UK£5.99, US & Eur £6.99, rest of world £7.99
DVD - UK £6.99, US&Eur £7.99, rest of world £8.99.

LINUX

www.linuxformat.co.uk **FORMAT**

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-1.0.1.i386.rpm – This is probably a binary rpm, designed to run on x86 systems.

Someap-1.0.1.i386.deb – The same, but a debian package.

Someap-1.0.1.tar.gz – This is usually source code.

Someap-1.0.1.tgz – Same as the above, tgz is abbreviated form of tar.gz

Someap-1.0.1.tar.bz2 – Same, but uses bzip2 compression instead of zip

Someap-1.0.1.src.rpm – This is also source code, but supplied as an rpm to make it easier to install

Someap-1.0.1.i386.RH7.RPM – A binary, x86 RPM designed specifically for Red Hat Linux

Someap-1.0.1.ppc.Suse7.rpm – A binary RPM designed specifically for SuSE7.x PPC Linux.

Someap-devel-1.0.1.i386.rpm – A development version.

Installing from tarballs

A tar ball is a two stage archive. First the files are archived into a single file with **tar** and then compressed with **Gzip** or **Bzip2**. To unpack, **cd** to the directory you want to unpack it, usually your home directory and type one of the following two lines:

```
tar xzvf /mnt/cdrom/Desktop/progname/progname-2.1.0.tgz
```

```
tar xvf -bzip2 /mnt/cdrom/Desktop/progname/progname-2.1.0.tar.bz2
```

Use the first for Gzipped files, those ending in .tar.gz or .tgz, and the second for Bzipped files, ending in .tar.bz2 or .tbz2. Naturally, you change the paths to suit the location and name of the archive. and replace /mnt/cdrom with whatever is applicable to your system (eg /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 CDs

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 information. [LXF](#)

Coverdisc



Neil Bothwick has tracked down twice the usual amount of software to shoehorn into this month's **DOUBLE-SIDED Linux Format DVD**.

This month's DVD contains a double bonanza of software. It is double sided, so after you have explored the usual selection of software, flip the disc over and start all over again! With 8GB of space to play with, we have been able to bring you more larger packages and distros yet still include hundreds of the smaller, but no less useful, utilities, applications and games. The distributions on this month's DVD include Lycoris, intended to give an easy introduction to the world of Linux, Knoppix, the popular "live CD" distribution, and Yellow Dog for PowerPC users.

Even though we have allocated another couple of pages to highlight

software on the DVD, this is only a drop in the ocean. Load the index.html file on each side on the DVD for a complete listing of what is on the DVDs this month, complete with descriptions, screenshots and links to the programs' home pages.

KDE 3.1 is in the Desktop directory of the first side of the DVD. Everything written on the CD pages about KDE 3.1 applies here too, with the exception of the locale files. The DVD version has all the internationalisation files included.

DISTROS LYCORIS

Formerly known as Redmond Linux, Lycoris is intended to be an easy to use, and easy to install, Linux distribution. This is not for the power

user or server admin, it is aimed squarely at the desktop user who just wants to get on with using it without worrying about how it works. There are three CD images on the DVD, but you only need the first one to install Lycoris. See the pages at the end of this section for information on creating CDs from the ISO images on the disc.

Installing Lycoris is intended to be as easy as possible, and it really is easy. Boot from the first CD to start the installer. At the splash screen, select "Install Desktop/LX". The screen goes blank for a while when the installer is loading. This can be a little disconcerting, but as long as the access light on your CD drive is active, it's nothing to worry about. After the usual promotional screen and licence agreement, you select your mouse,

keyboard and video card. It picked up my mouse and keyboard with no problems, but got the graphics card, a GeForce 2, completely wrong. If this happens to you, select the correct one and continue. When you select your monitor on the next screen, make sure you test it. If you don't and the video mode is inappropriate, you may not be able to get at the desktop.

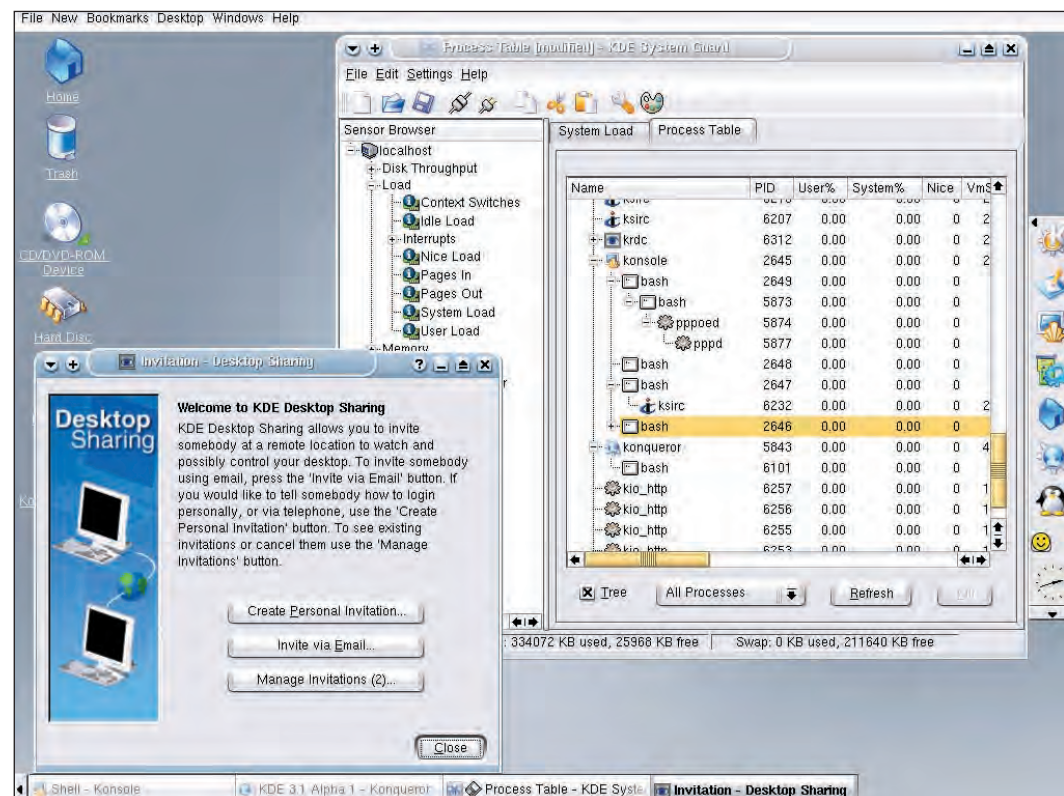
Now we come to what is the trickiest part of most distro installs, partitioning the hard drive. Unless you have empty space on your drive, or want to use the whole hard disk, you will need to select at least two partitions, / and **swap**. **swap** should be one to two times the amount of memory in your machine, / will contain everything else. You can either select existing partitions and reformat them, »



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

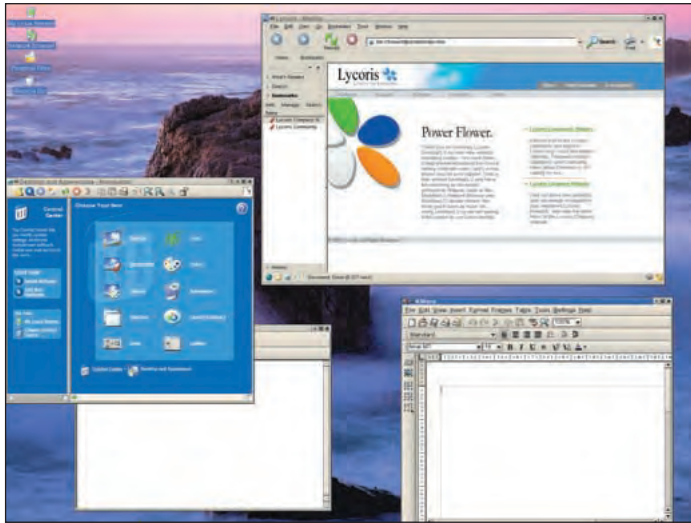
IMPORTANT NOTICE

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.



KDE's new desktop sharing lets you give other people or computers controlled access to your desktop.

CoverdiscDVD



Lycoris is intended to be an easy to install and use introduction to Linux for those users happy operating within a GUI most of the time.

or create new ones with the Repartition option. The next screen lets you format these partitions, which you will need to do before anything can be installed on them.

The next stage is creating accounts for normal usage and system administration. Type in your details, pick a password and press "Add user". Then give a password for the system administrator and press Next. You can add other users here too, or you can add them later. The next screens let you set up networking, a modem and a printer. A nice touch is that the installer is installing packages while you are giving this information, if you finish before it does, it gives you a game of patience to play while you wait. Finally, you have the option of creating a rescue disk, which you should so.

That's it! Click on Finish and your new Lycoris desktop will load. Press the menu button at the bottom left of the screen and start exploring the options. Make sure you log in as your normal user, *not* the system administrator. This will let you explore and experiment with the various programs without allowing you to do anything that could damage your new installation. When you have finished, for now, right-click on the desktop and select Logout from the menu.

DISTROS YELLOWDOG

Most machines running Linux use X86 processors, so it is inevitable that our cover discs concentrate most on packages for i386 and above. Most packages come with source code, so

users of other architectures can still compile them to run on their machines, and we do have a few PPC RPMs when they are available, but this is not the case with distros. So this month we have a Linux distribution for PowerPC users. Yellow Dog is based on Red Hat Linux and comes as a single ISO image which you should burn to a CD-R. Then boot your computer from this disc and follow the installation process. If you have problems booting from the CD, the support section of the Yellow Dog web site at www.yellowdoglinux.com/support/ contains plenty of information on how to install on the various models of Mac.

DISTROS KNOPPIX

We had Knoppix on the DVD a few months ago, as a bootable disc. It has been updated several times since then so we thought it was time to bring you the newest version, especially as the original inclusion of it proved so popular. This time it is supplied as an ISO image, ready for burning to CDR. Knoppix is a "live CD" distribution. That is, when you boot from the CD, you don't get an installer. Instead, it boots straight into a working Linux environment, based on Debian, complete with KDE, GNOME, Open Office and loads more.

There are two main reasons for using Knoppix. You can run it on any i386+ machine no matter what is installed on the hard drive. This is good for showing Linux to friends or colleagues, or just for giving yourself access to your favourite OS on

another computer. The other main use for it is as a rescue disc. Forget seeing how much you can cram onto a floppy disk, there is something like 1.7GB of software compressed onto this disc, giving you all the recovery tools you need, and then some. If you do not want to load up X and KDE for rescue work, type **knoppix 2** at the boot: prompt to boot into a console. Press F2 at the boot: prompt to see the other options.

Even though the distribution boots from CD, it is possible to save your configuration and home files. Put a blank floppy disk in the drive and select KNOPPIX->Configure->Save KNOPPIX configuration from the KDE menu. The next time you boot, simply type **knoppix floppyconfig** at the boot: prompt and your settings will be reloaded from the floppy disk.

One of the additions since the last time we covered Knoppix is that it can now be installed onto your hard drive with a minimum of effort. You will need free space for the root and swap partitions, then log into a console as root and type **knx-hdinstall**.

DESKTOP CYGWIN

The extra space of a double-sided DVD means that we are able to bring you some software to run on Windows. No, your eyes didn't deceive you, that really did say "software to run on Windows". However, most of it is not Windows software. Cygwin is a comprehensive Unix/Linux environment that runs on Windows. It means that even if you are unable to install Linux on a computer, you can still have access to a large range of

Linux and Unix software. Cygwin achieves this magic with a .DLL file that provides an emulation layer.

To install Cygwin, double-click the setup.exe file in the Cygwin directory. Select the bottom option, Install from local directory, and follow the prompts. When asked for a directory to save downloaded files, select the Cygwin directory on the DVD (this should be the default). Yes, it does seem odd, but the installer is mainly intended for an Internet installation, where the files are first downloaded to a local directory and then installed from there. We have done the downloading for you, but the installer still refers to the directory containing these files as a save directory. don't worry, it will work. Once the installation is complete, double-click the Cygwin icon on the Windows desktop to open a bash shell. Welcome home!


Once you have Cygwin installed and working, there is plenty of information on what you can do with it on the Cygwin web site at <http://cygwin.com/>

DESKTOP XNORTHERNCAPTAIN

KDE has an extremely capable file manager in *Konqueror*, improved for version 3.1. However, sometimes a simple two-pane file manager is a quickest way to get the job done when you have a lot of files to sort through. *X Northern Captain* (XNC) is such a program, and a good example of the breed. XNC allows you to do much more than the standard copy, move, rename and delete operations. It is able to handle FTP sites and archives in a similar manner to local



Knoppix, back by popular demand, and this time it is hard drive installable.



Recursion in action! Cygwin running a bash shell on Windows 98, which is running in VMWare on Mandrake 9.1.

filesystems. The menus and file associations are all configurable by the user, so you can customise this file manager to work in exactly the way you want, rather than having to adapt your way of working to fit in with the software.

XNC comes with its own file viewer and editor, although you can reconfigure the program to use whatever viewers and editors you feel most comfortable with. Bookmarks let you jump quickly to your most used directories without navigating the filesystem and hotkeys give fast access to the most commonly used functions without the necessity of using the menus. **XNC** is supplied as binary and source RPMs as well as binary and source tarballs, so most installation preferences are covered.

GAMES SENKEN

How often do you drive through your home town or city, cursing whoever planned the roads, put the shops the other side of a busy bridge and generally made life far more difficult for you than it should be? Now you have the opportunity to prove how you could do a much better job. **Senken** is a city simulation game – in much the same vein as the popular Windows-based *SimCity*. You must buy the land, build the infrastructure, balance the books and convince people to move in. You can play to achieve targets and construct an effective conurbation, or just fool around to see if you can screw thing up as well as your local planners.

SERVER DANSGUARDIAN

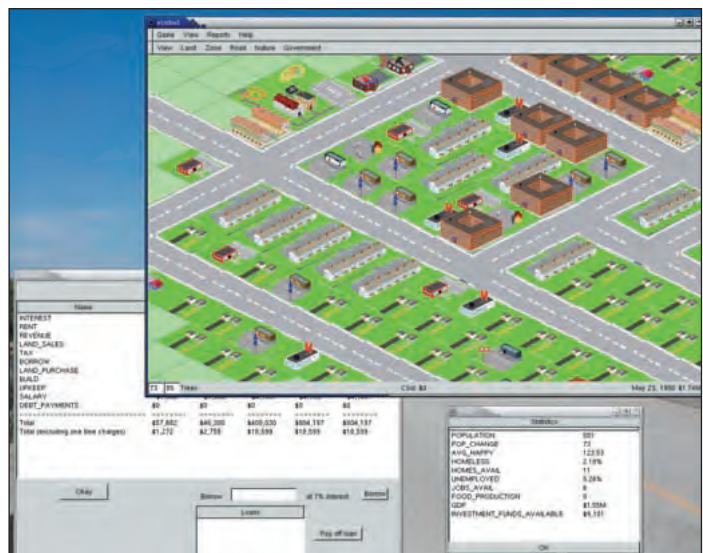
If you have children who access the Internet, you are probably concerned about the kind of content they could stumble across while browsing the web. **DansGuardian** is a web filter that is able to block access to sites using a variety of methods.

DansGuardian does not filter simply by blocking URLs from a blacklist, it parses the content of pages and blocks them based on various phrase-matching algorithms. It is up to you how strictly you control access. The default settings are suitable for children of primary school age, but you can set it to be as strict or as lax as you like. It can also filter on URLs and PICS ratings for sites, but its real strength is that it is able to filter the actual content of web pages, rather than where they are coming from.


DansGuardian works in conjunction with the *Squid* proxy server, which is almost certainly on your distribution CDs. Because it works through a proxy server, it is able to protect all machines on a network, not just the one it is installed on. The other machines don't even need to be running Linux. If you do have Windows machines on the network, you should also look in the *DansGuardianVirusScanner* directory. This adds the ability to scan all content for viruses before it reaches the computer downloading it.

OFFICE ABIWORD

OpenOffice.org seems to get all the publicity in the Open Source word



Design and build your own city, then convince people they want to live there. See if you can do better than your local council with *Senken*.



X Northern Captain is a file manager using the traditional two-pane format, but with some interesting and useful features.

processor stakes, which is hardly surprising as it has such a high profile company behind it. However, there have been other Open Source word processors around for a lot longer. KDE's *KWord* is one, and *AbiWord* is another. It started life as part of *AbiSource*, a project to create a cross-platform office suite. *AbiSource* is no more but the main component, the word processor, lives on. One of the main considerations of *AbiWord* is that it should remain lean while still having a good set of features. It should appeal if you find *OpenOffice.org* heavy going on your hardware.

If you are an *OpenOffice.org* fan, you'll be pleased to know that they have released an update. This is a bug fix update. It doesn't add new features but is reported to be slightly faster,

with improved printing. Naturally, we are not telling you this just to tease you, the latest version is on the DVD.

DEVELOPMENT ECLIPSE

Eclipse is an open extensible Integrated Development Environment (IDE) for anything and nothing in particular. Instead of being tied to a specific programming languages, **Eclipse** uses tool plug-ins that make it work with a wide variety of languages, from C and C++ to Java and PHP. There are two versions on the DVD, differing only in their use of GTK or Motif, choose whichever look you prefer. The archive itself contains a number of plugins, there is also a PHP plugin in a separate directory on the DVD. **LXF**

CoverdiscDVD

DVD CONTENTS AT A GLANCE

Side 1

Magazine

Python A high-level scripting language
RoundUp All the programs covered in this month's Round Up feature
The GIMP The GNU Image Manipulation Program

Desktop

Batchelor Home resource planning software
CDDetectAndExecute A CDROM disc type detector
Cygwin A UNIX-like environment for Win32
GenDoc An XML editor written in Java 2
GNOMECommander A GNOME based filemanager
HistoricalEventMarkup An XML/Java markup and display of historical events
KDE3.1 A powerful graphical desktop environment
Kile An integrated LaTeX development environment for KDE 3
KPreg A helper application for completing Web forms
MicrosoftWord2002Unmunger Removes HTML cruft from Word2002 documents
MultiSync A calendar synchronization program,
PHPSambaExplorer A PHP-based Web interface to smbclient
Qastrocam A Webcam and telescope program for astronomy purposes
SmartMonTools Utilities to monitor S.M.A.R.T. disks and devices
Stellarium A 3D astronomical sky renderer
Txt2Tags A generic text converter
Wine Emulator of the Windows 3.x and Win32 APIs
XFileExplorer A file manager for the X Window System

Development

Bugzilla A bugtracking system project hosted and used by mozilla.org
CVSToys A CVS enhancement that provides commit notifications
DocBook DocBook stylesheets, schemas, documentation, and resources
DocBookDoclet Create DocBook code from Java source documentation or HTML
Eclipse A universal tool platform
FormatJavaForEclipse Format multiple Java source files in one go
GOBOEiffel Open source and portable Eiffel tools and libraries
Houston Status and logging toolkit
PHPEclipse PHP support for the Eclipse IDE Framework
PyChecker A tool for finding common bugs in Python source code
Python-dpkg Tools to handle dpkg files and data types
SimpleCExpatsWrapper Expat wrapper around that provides an object model
Soya3D A high level 3D engine for Python
SybaseModuleForPython A Python DB-API-compliant interface to Sybase
wxPython A Python extension module for wxWindows
wxWindowsGTK GTK port of the cross-platform wxWindows
XMLSecurityLibrary Implementation of XML Signature and XML Encryption standards

Internet

Amaya The browser/authoring tool of the W3C
Fetchmail Full-featured, robust, well-documented mail retrieval
FXservIRCServices Web searching, SMS, file serving and proxy detection
InternetRelayJabber A Jabber client for IRC
Junkie A GTK 2 FTP client with many features
ProxIRC Enables simultaneous use of IRC and MSN using any IRC client
Screen-Scraper A tool for extracting data from Web sites
Sylpheed A GTK+ based user-friendly email client
VaporMail A J2EE enterprise personal information management system
VOCP Complete voice messaging system
WhoisDomainCheckup A Web-based domain name checker written in PHP

Office

AbiWord A fully-featured word processor
BemDB Group health insurance database
GnuCash A program to keep track of your finances
Gnumeric A spreadsheet that's part of GNOME
MrProject A project management application
OpenOffice.org An Open Source version of StarOffice
Verlof A Web application for managing leave information

Side 2

Distros

Knoppix A bootable CD with a collection of GNU/Linux software
Lycoris A Linux distribution aimed at ease of use.

PhayouneFirewall A CD-ROM based firewall distribution
PXESLinuxThinClient Micro distro for converting PCs into thin clients
TopologiLinux A Linux distribution designed to coexist with Windows
WISP-Dist A modular embedded Linux distribution for wireless routers
YellowDog A complete port of Red Hat's RPM-based OS to PPC

Games

Celestia 3D space simulation
Epiar A space adventure/combat game
GamesForGnome2 A collection of games for the GNOME2 desktop
K15 The KDE version of the fifteen game
KingOfTheHill Multiplayer artillery game in the scorched earth tradition
Kugellines A board game where you create lines of balls
Palito A real-time strategy game
Poopmup A shoot-em-up, with a twist
ScummVM Interpreter for LucasArts and AdventureSoft adventures
Taruli The classic game of concentration

Graphics

AstroMosaic Helps in the construction of a big mosaic from small images
Evolvotron A "generative art" image evolver
KDEDivXSubtitlesEditor DivX subtitles editor for KDE 3.x
MPlayer A movie player for Linux
SWFTools Read, write, generate and modify Flash files

Mobile

at76c503LinuxDriver A Linux driver for at76c503-based wlan USB adapters
KDEPocketPCContactsImport Import WinCE/PocketPC Contacts into a KDE Address Book
PicoGUI A complete GUI system for embedded systems

Server

BannerAdRotationProgram Sell out and manage banner ads with Perl and SSI
Checky An interface to online Validators and Checkers
DansGuardian A Web filter with multiple page content analysis mechanisms
DansGuardianVirusScanner Adds anti-virus scanning support to DansGuardian
Harvest A Web-based document search system
JBoss An Enterprise JavaBeans application server
Jigsaw W3C's leading-edge Web server platform
LazarusNewsreel A Web site news publishing application
mod_variety Apache module to serve a random file
phpDealerLocator Keep track of your dealers
phpMailAdmin A Web-based admin system for mail accounts
PhpWiki A WikiWikiWeb clone written in PHP
PowerMail Distributed system for receiving and storing mail
PPMy Web-based administration utility for ProFTPD
PVote A PHP voting system
QuantaPlus A web editor for KDE supporting HTML, PHP, and more
Sporum A better web-based discussion board software
Webtop Web-based email and PIM
Zorbstats Stats for your Web site using PHP and MySQL

Sound

BurnCDDA A tool for creating audio CDs
GStreamer Streaming-media framework
Manauton Records sound to disk in manual or autonomous mode
OpenBeatBox The Virtual Drum Machine
UnattendedBroadcasting A daemon for running the operations of a radio station
XSIDPLAY A C64 music player and SID sound chip emulator

System

AutoPKG A Slackware update / upgrade / install tool
CommonUNIXPrintingSystem An Internet printing system for Unix
Firestarter Firewall creation/monitoring tool
FirewallBuilder Multi-platform firewall management software
JMapPortScanner A port scanner for individual hosts or network subnets
MACChanger Manipulate MAC addresses of network interfaces
MindILinux A Linux mini-distro derived from your current system
MondoRescue Generates disaster recovery CDs or tapes
OpenLDAP LDAP suite of applications and development tools
OpenSSL Toolkit for SSL and Transport Layer Security
Tcpreplay Replays captured network traffic back onto the wire
X10UniversalDeviceDriver A device driver for X10 home automation for script access
YetAnotherNVConfigurator GUI configuration tool for the Linux nVidia driver set

User Groups

LUGs worldwide are full of members keen to help with your problems, discuss ideas, and generally natter about all things Linux. You can find lots more information online at: www.lug.org.uk

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

4 OXFORD

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

5 KENT

URL www.kent.lug.org.uk
Contact Kevin Groves

REVISED
DETAILS

6 BRIGHTON

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

7 WORCESTERSHIRE

URL www.worcs.lug.org.uk

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.wolveslug.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://glug.linux.co.uk/>
Contact John Southern

20 SURREY

URL www.surrey.lug.org.uk
Contact Jay Bennie

21 CAMBRIDGE

URL www.cam-lug.org.uk

22 DEVON & CORNWALL

URL www.dclug.org.uk
Contact Simon Waters

23 FALKIRK

URL www.falkirk.lug.org.uk

24 MANCHESTER

URL www.man-lug.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.shefflug.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 BERKSHIRE & 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.sderbylug.org.uk
Contact Dominic Knight

36 BELFAST (BLUG)

URL www.belfastlinux.cx
Email russell@belfastlinux.org

37 WILTSHIRE

URL www.wiltshire.lug.org.uk
Contact Jason Rudgard

38 SOUTH LONDON

URL www.sl.lug.org.uk
Email edo@perceptiondm.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 <http://midlandslug.port5.com/>
Contact Pete Thompson

42 CUMBRIA

URL www.cumbria.lug.org.uk
Contact Jamie Dainton

43 DORSET

URL www.dorset.lug.org.uk
Contact John Robinson

REVISED
DETAILS

44 SHROPSHIRE

URL www.shropshire.lug.org.uk
Email shropshire@lug.org.uk

45 SOUTH WEST

URL www.southwest.lug.org.uk
Email southwest@lug.org.uk

REVISED
DETAILS

46 SOUTH WALES

URL www.swlug.org.uk

47 NORTH LONDON

URL www.kemputing.net/lug/anlug-aims.html
Email jason@voyagercomputers.co.uk

48 MALVERN

URL www.malvern.lug.org.uk
Contact Greg Wright

49 HUDDERSFIELD

URL www.hud.lug.org.uk
Contact Dave Naylor – knocker@caramboo.com

REVISED
DETAILS

50 NOTTINGHAM

URL www.nottingham.lug.org.uk

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

57 BLACKBURN

Email matt@consultmatt.co.uk

58 YORK

URL www.york.lug.org.uk

59 Lincs

URL www.lincs.lug.org.uk

**60 HULL**URL www.hull.lug.org.uk**61 WALTON-ON-THAMES**Contact William Mutch
Email rael@freeuk.com**62 GLOUCS & COTSWOLDS**URL www.gloucs.lug.org.uk**REVISED
DETAILS****63 WEST OF SCOTLAND**URL www.wos.lug.org.uk**64 SOUTH STAFFORDSHIRE**URL www.staffs.lug.org.uk**65 MANSFIELD**URL www.mansfield.lug.org.uk**66 BORDERS**URL www.linux.bordernet.co.uk**67 BIRMINGHAM**URL www.sb.lug.org.uk**68 COVENTRY**Email info@coventry.lug.org.uk**69 NEWARK & LINCOLN**URL www.newlinc.lug.org.uk**70 BEDFORDSHIRE**URL www.beds.lug.org.uk**71 LINCOLN**URL www.lincoln.lug.org.uk**72 LOUGHBOROUGH**URL www.loughborough.lug.org.uk**73 EXETER UNIVERSITY**Email N.J.Murison@exeter.ac.uk**74 SUNDERLAND**Email thomas.croucher@sunderland.ac.uk**75 EAST YORKSHIRE**Email sharkonline@whatemail.com**76 CLEVELAND OPEN SOURCE GROUP**Email openlug@digitalmedia.co.uk**77 BEVERLEY**Email vladimir_lukyanov@hotmail.com**78 DUNDEE & TAYSIDE**URL www.dundee.lug.org.uk**79 SUSSEX**URL www.phpworld.co.uk/~swlug**80 WIGAN & ST HELENS**Email paulf.johnson@ukonline.co.uk**81 BRIXTON**URL www.communitytechnology.org.uk/~linuxhome**82 ST.ANDREWS, FIFE**URL www.standrews.lug.org.ukEmail stuart@nx14.com**83 NUNEATON**URL www.nuneaton.lug.org.uk**84 ISLE OF MAN**URL www.iom.lug.org.ukEmail helix@manx.net**85 AYLESBURY**URL www.aylesbury.lug.org.ukEmail drbond@educational-computing.co.uk**86 LANCASHIRE**URL www.lancasterlug.org.uk**87 EAST LONDON**URL www.eastlondon.lug.org.uk

Contact Jonathan Spriggs

NEW DETAILS:**88 ORMSKIRK**Email rob@northwestlinux.co.uk**89 HEREFORD**URL www.hereford.lug.org.uk/Email rbjh@good-news.fsnet.co.uk**90 EAST HERTS**Email madtom1999@yahoo.com**91 SWINDON**Email nick.trueman@ntlworld.com**92 MENAI**URL www.menai.lug.org.uk**93 ABERDEEN**Email mark@aricc.net**94 SHETLAND**URL www.shetland.lug.org.ukEmail c_s_butler@yahoo.com**NORTH STAFFORDSHIRE****Meeting** on 18th February to kick-start a North Staffs LUG. Details on <http://titanic.co.uk/staffslug>**OTHER GROUPS:****B-LUX**URL www.b-lux.org/Email blux-admin@lineone.net**FUENGIROLA GNU/LUG**URL www.fuengirola.lug.org.uk/Email fuengirola@lug.org.uk**SCHOOLS**URL www.schools.lug.org.uk/**TR LINUX – TURKISH LINUX**URL www.trlinux.lug.org.uk/Email c_s_butler@yahoo.com**YOUNG LINUX**URL www.young.lug.org.uk/

LUG OF THE MONTH

Scottish Linux Users Group

Ben Thorpe writes: we're currently awaiting the appointment of a new 'head honcho' after the retirement of our last leader. Though our name relates to the whole of Scotland, our meetings are in the Glasgow area, though we always like to hear from any Caledonian Linux users! We meet on the last Thursday of every month, beginning in the CIS Staff common room (L11.09) Livingstone Tower, Strathclyde University, usually with a Q & A session and a talk on a Linux-

related subject, and then continuing at 9pm in the Counting House off George Square for less structured debate. Maps are available on the website. Membership is varied in both age and ability. January's talk was on Content Management Systems, with a particular tip-of-the-hat to Zope. February's talk will be on regular expressions.

SLUG was established in 1998, and attributes its longevity to never asking for money, and never forming a

committee :) I've only been active on the mailing list for a couple of years, and going to the meetings for a year or so, but have found it to be a large, yet welcoming, group of geeks.

Visit www.scotlug.org.uk. IRC is currently **#scotlug** on **OpenProjects.net**. There is also an active mailing list, which you are welcome to subscribe to at <http://mailman.lug.org.uk/mailman/listinfo/scottish>



Worldwide Linux User Groups

Free Software users across the globe

Africa

EGYPT

URL www.linux-egypt.org

Contact Hesham Bahram

GAUTENG, SOUTH AFRICA

URL www.glug.org.zaEmail glugmin@revolution.org.za

Australia

ADELAIDE

URL www.linuxsa.org.auEmail mtippet@anu.edu.au

ALICE SPRINGS

URL www.aslug.org.au

MELBOURNE, VICTORIA

URL www.luv.asn.auContact luv-committee@luvasn.au

PERTH

URL <http://plug.linux.org.au/>

SYDNEY

URL www.slug.org.au

Europe

AUVERGNE

URL www.linux-arverne.orgEmail Cyril.Hansen@wanadoo.fr

COSTA DEL SOL (English speaking)

URL www.fuengirola.lug.org.uk

DENMARK

Alssund www.alsslug.dkEsbjerg www.eslug.dkFyns www.flug.dkMidt-og Vestjylland www.mvjlug.dkNordjylland www.njlug.dkSkåne Sjælland www.sslug.dkTrekantsområdet www.tlug.dkVest-fyn www.haarby-net.dk/vflugÅrhus www.aalug.dk

EIRE

URL www.linux.ieEmail root@linux.ieURL www.dilu.orgEmail glossary@dilu.org

GOTHENBURG

<http://nain.oso.chalmers.se/LUGG/>

Middle East

ISRAEL

URL www.iglu.org.il/IGLU/Email webmaster@iglu.org.il

PALESTINE

URL www.lugps.orgEmail isam@planet.edu

North America

ALASKA

URL www.aklug.orgEmail deem@wdm.com

ALBERTA

URL <http://calgary.linux.ca/>

BATON ROUGE

URL www.brlug.netEmail dpuryear@usa.net

BAY AREA

URL www.balug.orgEmail aftyde@balug.org

CLARKSVILLE, TN

URL www.clug.orgEmail tux@clug.org

DENVER

URL <http://clue.denver.co.us/>

FLORIDA

URL www.flux.org

LOS ANGELES

URL www.lalugs.orgEmail dank@alumni.caltech.edu

NORTH COLORADO

Email nclug@nclug.org

OTTAWA CANADA

URL www.oclug.on.ca

TAMPA

URL www.suncoastlug.orgEmail president@suncoastlug.org

UHACC Normal, IL

URL www.uhacc.orgEmail lug@uhacc.org

VIRGINIA TECH

URL www.vtlug.orgEmail nega@vt.edu

South America

BUENOS AIRES

Email dcoletti@impost.com.ar

CHILE

URL www.linux-chile.org

MONTEVIDEO

URL www.linux.org.uy

PARAGUAY/ ASUNCION

Email rolgiati@conexion.com.py

SAO PAULO

URL <http://gul.ime.usp.br/>Email gul@ime.usp.br

Spreading the word

BBLUG's Install Day expected a trickle, and coped with a torrent. Amias Channer took his umbrella and wellies.

Imagine a scene something like

Lourdes, where the faithful masses gather to pray for divine intervention to help their afflicted friends. Now transfer that image to Bath on a cold January Saturday. The faithful become Linux geeks (and wannabe geeks) from all over Somerset, Gloucestershire and Wiltshire, and the gurus were anyone who might have some useful information. The holy-men imparting wisdom included employees from the likes of HP, IBM and many local contractors, all giving up their time for no financial reward – thanks folks, we really appreciate your knowledge and dedication to Linux! While X86 Linux issues dominated the event, there was a smattering of sparc, mac, mini-ix and even hppa hardware being admired, traded and upgraded. Other more frivolous activities included a spectacular home-brew joystick (with real arcade buttons!) connected to a laptop running classic arcade games via MAME, some scary soldered

serial port hackery for Palm Pilots, and the old favourite – the 'put faces to names that I've met online but never in real life' game.

The whole event was buzzing with the cheerful and helpful spirit that characterises Linux user groups. Everyone was agreed that another event should happen soon, but at a bigger location. Although the bandwidth (provided by Shevek, the organizer) coped admirably (even with the Gentoo install in the garden hoovering up all it could manage), it was clear from the turnout (60 or so) that organisation on a larger scale is required next time. It is at events like these that it's possible to get a feel for how Linux came about. By providing a place for Linux users to socialise, share knowledge, work directly on hardware, swap components or even just admire the results of each others' hard work, the enthusiasm that drives the development of Linux is thoroughly recharged all without any company being officially involved. Support your local Linux 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**Production Editor** Matthew Nailon
matthewnailon@futurenet.co.uk**Editorial Contributors**

Jono Bacon, Neil Bothwick, Chris Brown, David Cartwright, Andy Channelle, Amlas Channer, David Coulson, Chris Denton, Richard Drummond, Hoyt Duff, Alan Ferguson, Marco Fioretti, David Harris, Paul Hudson, Biagio Lucini, Patrick K O'Brien, Richard Smedley, Ben Thorpe, Bill Von Hagen, Tom Wilkinson

COVER CD PRODUCTION**CD Editor** Neil Bothwick
Senior CD Editor Kate Hadley**ART CONTRIBUTORS****Photography** James Wilson, Louise Parker, Getty Images, Superstock, Corbis**Illustration** Paul Bateman, Chris Winn, Shane Collinge, Michael J Hammel, Dylan Teague (2000 AD/Rebellion)**ADVERTISING SALES****Portfolio advertising manager** Matt Stanley 01225 442244
matt.stanley@futurenet.co.uk**Portfolio deputy ad manager** Diane Clydesdale**Business development manager** Damian Hughes**Account managers** Zoe Drew, Giles Crosthwaite Scott,**Ad sales executives** Jo Doyle, Stephen Hall**MARKETING AND PROMOTIONS****Marketing Manager** Fiona Tully
Subscriptions Executive Marie Spicer**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** +44 (0)1225 442244.**Overseas Licences****International Licensing Director** Simon Wearsimon.wear@futurenet.co.uk **Tel** +44(0)1225 822798**Fax** +44 (0)1225 788105 **Cell phone** +44(0)777 5641493**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 96**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 TPL
© Future Publishing Ltd 2003

LINUX is a trademark of Linus Torvalds, GNU/Linux is abbreviated to Linux throughout for brevity.

All other trademarks are the property of their respective owners

Future Publishing Ltd. is part of The Future Network plc.

The Future Network produces carefully targeted specialist magazines and websites for groups of people who share a passion. We aim to satisfy their passion by creating titles that offer superb value for money, trustworthy information, multiple ways to save time and money, and are a pleasure to read or visit. Today we publish more than 80 magazines and over 30 magazine websites and networks from offices in five countries. The company also licenses 32 of its titles resulting in over 60 local editions in a further 23 countries.

The Future Network plc is a public company quoted on the London Stock Exchange (symbol: FNET).

Non-executive Chairman: Roger Parry

Chief Executive: Greg Ingham

Chief Operating Officer & Managing Director, UK: Colin Morrison

Group Finance Director: John Bowman

Tel +44 (0)1225 442244

BATH • LONDON • MILAN •

NEW YORK • PARIS •

SAN FRANCISCO • WROCLAW

www.thefuturenetwork.plc.uk



Media with passion

NEXT MONTH

Issue 39 on sale Wednesday 26 March

FILE MANAGERS Head-to-head

The biggest ever roundup of Linux filemanagers – are you using the one that's best for your needs? Will Konqueror win out, or will something else take command? The best-known filemanagers and some unknown quantities battle it out to be King Of The Desktop.

PLUS:

Serial networking

Ethernet networking may often be the preferred option, but we tackle the concepts, whys and wherefores of networking over standard serial links.

High Street Linux on test

Perhaps installing Linux will become a thing of the past as PCs pre-installed with everyone's favourite OS start to ship. We take a close look at Evesham Micro's eEscape Linux desktop.

And get the latest issue of Linux Pro!

Featuring more on issues concerning professional Linux users, including: firewalls, storage, Java servlets and a look at the most powerful Linux systems in the world!

DON'T MISS YOUR COPY OF LINUX FORMAT!

Get it delivered to your door every month – subscribe on page 96

The exact contents of all future issues are subject to change

LINUXPRO

FROM THE MAKERS OF LINUX FORM

MARCH 2003

ONLY AVAILABLE WITH
LINUX
FORMAT

**NEW
BIGGER
SIZE!**

PLUS

Apache Rewriting

Save time, space
and effort on your
virtual hosts

IP Tables in depth

Building a firewall
from scratch is often
the best solution

Pocket-sized Linux

Exploring the new
Montavista
Consumer
Electronics Edition

Tomcat Java

Deploying Java
servlets and JSP the
Open Source way

To Protect and serve

Construct a backup strategy that works
for you, your budget and your data **PAGE 20**

Welcome

Twenty-four pages of real-world Linux for IT professionals

It's all-change again in the world of *Linux Pro*. This issue we have no fewer than three new short series and a new storage section for you. Storage is becoming a bigger and bigger topic in the world of IT, as new technology, rich media and ballooning filesizes stir the mix. We will be covering all aspects of storage technology, from configuring RAID to setting up a SAN system, with plenty of reviews and comparisons of the latest drives, tape devices and other storage hardware. This issue, as you may have guessed from the cover, we are looking at the world of backups, and more importantly, developing a strategy to ensure that your backup solution protects you properly by matching your needs.

Firewalling also falls into the security/protection theme. Many use off-the-shelf firewall devices or GUI-driven firewall tools, but you may be missing out on some of the advanced features available with the 2.4 kernel's built-in *iptables* code. In this first of a two-part series, we look at the bare bones of *iptables*, and how you can build a robust firewall with just a few simple rules.

Running Java servlets the open source way is introduced in another new series this issue, focussing on the *Tomcat* software project. It's more than just a servlet container as you will learn, and offers a degree of flexibility not found in all of the commercial offerings.

Do you really need a Grid? Probably not, claims Alan Ferguson of the Anix group, in an opinion piece on page 6. Grid computing may well be the most hyped technology, but what can businesses learn from it, and who will actually be able to make use of it?

If you're more interested in smaller-scale Linux implementations, you might want to check out Montavista's new Consumer Electronics Edition over the page, which promises a deep and feature rich environment for embedded development. I hope you enjoy this issue – your comments, as always, are welcome at the address below.

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



“Storage is becoming a bigger and bigger topic in the world of IT as new technologies, rich media and ballooning filesizes stir the mix.”

Contacts

EDITORIAL

Editor
Nick Veitch nick.veitch@futurenet.co.uk
Art editor
Julian Jefferson julian.jefferson@futurenet.co.uk
Production editor
Matt Nailon
matthewnailon@futurenet.co.uk
Contributors
David Coulson, Alan Ferguson, Chris Brown, Ed Haynes

Photography
Superstock, Getty Images
Publisher **Dave Taylor**

ADVERTISING

Portfolio ad manager
Matt Stanley matt.stanley@futurenet.co.uk
Business development manager
Damian Hughes damian.hughes@futurenet.co.uk
Account managers
Zoe Drew, Giles Crosthwaite Scott
Ad sales executives
Jo Doyle, Stephen Hall

LinuxPro is a part of
LINUX
FORMAT
magazine

Published by
the future network
Media with passion

30 Monmouth Street, Bath BA1 2BW
Tel +44(0) 1225 442244

MARCH 2003

CONTENTS

LINUX IN YOUR POCKET
Montavista CEE Edition **p4**


GRID COMPUTING
Does your business need it? **p6**

TO PROTECT AND SERVE
A better backup strategy **p8**

APACHE REWRITING
Using the Rewrite module **p12**

TOMCAT
Java servlets and JSP **p16**

IPTABLES
Firewalling from scratch **p20**



8

20

LINUXPRO 3

Montavista Consumer Electronics Edition 3.0

Linux in your pocket

Montavista is one of the best known names in embedded Linux development, for many years providing an embedded solution based on a genuinely Open Source code, tuned and optimised to the specific needs of embedded developers. The newly announced Consumer Electronics version is a further refinement of this specialised Linux distro. Why the need for a specific version, and who is it targeted at?

"We're seeing that traditional consumer electronics, the kind that you turn on and off and are very fixed-function, is evolving into advanced consumer electronics, as well as new categories of devices emerging that are network connected, very complex devices with high performance hardware and sophisticated software." Says Scott Hendrick, senior product marketing manager for Montavista.

"A great example is from one of our customers – Sony. In November 2002 they launched the Cocoon channel server. Unfortunately they haven't released it yet in the US or Europe, but in Japan. In their booth at CES, they talked a lot about it. It is a strategic new product for Sony, powered by Montavista. It looks like a set-top box that would connect to

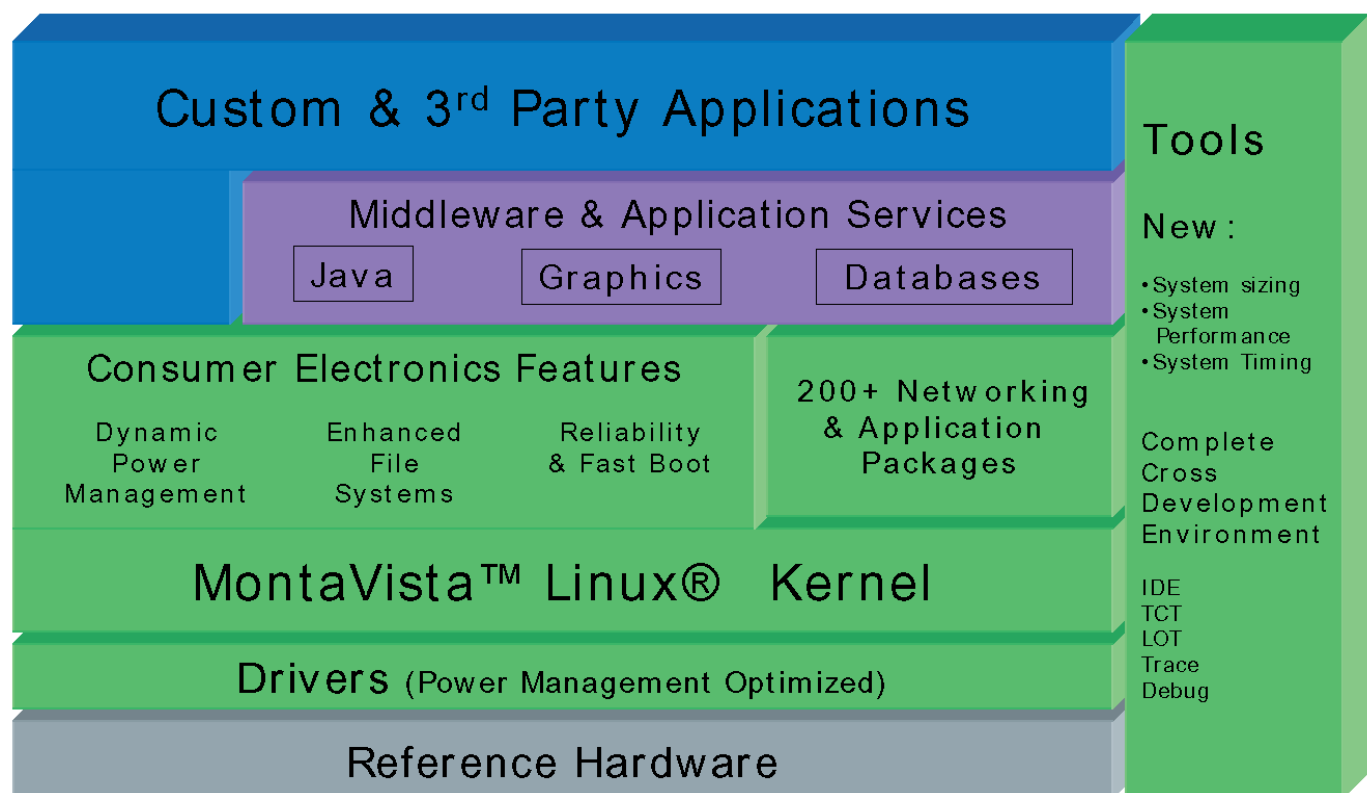
As the market for embedded Linux blooms, NICK VEITCH investigates a new release from Montavista aimed at consumer electronics.

How the Montavista CEE system stacks up.

your TV. It has a broadband connection and a lot of the PVR features. In addition to these features, it also has the ability to be accessed by mobile phone, so you can access it and get it to record something that you may be missing. More importantly, Sony have stated that this is a strategic platform for them and that over time they will be making different features and services available to the box. They will upgrade it over the network. In fact at CES they showed features that hadn't been available at the launch in Japan. So it's really a platform for service delivery."

Linux is very big in the embedded arena, and growing rapidly according to analysts such as IDC, Evans Data and others. The arguments for using a truly uniform and Open Source development platform are quite compelling. Not only can you draw on a wealth of available Open Source material, and the commercial middleware solutions, but developers can also use a stable and known development environment.

Linux also gives a degree of freedom, as Scott Hendrick explains: "Flexibility and control is really critical to our customers. Companies like Sony, NEC and others that want to be able to differentiate their products. They don't



want to be forced into a box or commoditized." These are all good cases for using Linux, but Montavista makes a case for its specific versions by very active development on key technologies of interest to the embedded market.

Apart from assuring the reliability of the kernel and the key peripheral drivers, they have also been responsible for many improvements to the standard kernel (the pre-emptable kernel patch for instance), and for implementing new features (eg O1 scheduling) which are still in the test kernel tree in a stable and useable environment.

Who's using it?

Over 100 consumer electronics companies are already customers of Montavista, including Sony, Panasonic and NEC. This includes a huge gamut of consumer devices ranging from set-top boxes and PVRs to PDAs, remote controls and VoIP telephones, and there are of course plenty more devices under development.

Big areas of growth for Linux include mobile phones (NEC have announced that they will be using Montavista Consumer Electronics for mobile phones), TV and Home devices and automotive telematics. This last category is expected to be a big area for Linux as companies try to reduce costs of advanced electronics in cars, both in the engine monitoring/management field and drivers' 'toys' like GPS, which currently can account for up to half the cost of the vehicle! All of these areas are specifically targetted by the new Consumer Electronics version of Montavista.

Making the grade

Consumer grade products have different expectations from Linux. For example, you might not mind your desktop computer taking a minute to boot up, but you wouldn't tolerate it from your mobile phone. Similarly you don't want your remote control to ask you whether you want to fsck the filesystem. As well as being reliable, a consumer grade OS has to be pretty much foolproof.

This means no kernel panics, good responsiveness, recovery from unexpected power loss and more. Montavista Consumer Edition implements all of these features and more. It also means that when things go wrong, the end user isn't going to be too inconvenienced.

"Our customers want to deliver a consumer electronics experience, not a PC experience. No blue screens of death – it's unacceptable to have to look for a paperclip to reboot your TV. We are very focussed on helping our customers to achieve," comments Hendrick.

One of the key technologies offered by the Montavista Consumer Electronic Edition is Xip, or 'execute in place' using the CRAMFS filesystem. This protects code, usually in EPROM or EEPROM, so that in normal use, the code can't be changed – it's read-only. As the code is executed direct from the read-only space, it can't be interfered with either. This means failsafe routines and critical software like the kernel, update routines and critical services (perhaps networking) can always be relied on, no matter what has happened to the rest of the filesystem.

"One important factor is the rich support for reference boards. Many of these systems – new chips for mobile phones, HDTV, telematics and so on – are highly integrated

CORE FEATURES

The horsepower under the bonnet of Montavista CEE Linux

- 2.4.20 Linux kernel
- Preemptible Kernel
- O(1) Scheduler

Leading GNU toolchains

- GCC 3.2
- GDB 5.2

Networking & Applications

- TCP/IP (IPv4 and IPv6),
- UDP/IP and PPP
- Apache, Thttpd, Samba, Sendmail, sshd, telnet, FTP, OpenSSL, OpenSSH, SNMP, mailx, rlogin, DHCP
- tcpdump, nfs-utils, snmputils

- iproute, routed, rtmon, Zebra (BGP/OSPF/RIP),

Language Support

- C, C++, Perl, python, shell scripts

Filesystems

- Ext2, Ext3, NFS, JFFS2, CRAMFS, RAMFS

MontaVista cross-development environment

- Target Configuration Tool
- Library Optimizer Tool
- Linux Trace Toolkit
- Kernel Debug Support (KGDB and Abatron)



The 10GB terrapin Mine, using existing Montavista Linux, is a powerful internet enabled personal storage device for images and audio tracks.

systems with many 'peripherals' included on the chip. Most support packages from most companies, including Montavista's professional version, usually only support a few of the key components and peripherals – Ethernet, framebuffer and so on. With our Consumer electronics version, we are going to be offering very rich support of all of the peripherals."

Typically on embedded applications, power management is an issue. The Montavista Consumer Electronics version aims to enable power management for all supported devices so they can be individually suspended and resumed – great for saving battery power.

Developers

Of course, as far as development is concerned, you can pretty much use standard Linux tools. GCC 3.2 tools are supported fully under this version of Montavista Linux, as is the GDB debugging tool. Additionally, Montavista make several other useful development tools available, including

"Linux provides a wealth of Open Source material, commercial middleware, and a stable development environment"



Everything from thin terminals to wristwatches are targets for embedded Linux.

a target configuration tool and a library optimiser. Various tuning and tracing tools are also available to help optimise the storage size and performance of your resultant embedded environment.

Initially the Consumer Electronics edition supports the TI OMPA 1510/5910 and IBM Power PC 405LP Arctic II boards, with more to follow. Expect to see a great horde of Linux enabled gadgets being released over the next year!

For further information on the Montavista Consumer Electronics Edition, and other Montavista products, please visit the Montavista website at www.mvista.com ■

GRID COMPUTING

Grid computing

Is it really worth all the hype?

Jet engine maker Pratt & Whitney may like it, NASA seems to think it's a good thing, and Boeing reckons it can now hardly do without it. Big names indeed supporting Grid computing, which along with Web services has become one of the hottest tech topics around. Its appeal is easy to understand: instead of buying expensive mainframes or supercomputers to process specific applications or jobs, use the otherwise idle processing power of dozens, hundreds or even thousands of existing PCs and servers to simultaneously process separate bits of data for any given task. The result is not only a cheaper hardware option, but a faster one as well. And in today's tight-fisted economic climate, cost-effectiveness remains very high on the IT investment priority list.

Not ones to miss such an opportunity, IBM, Sun and HP have for the past two and a half years been promoting Grid computing as the sea change that will drown all existing approaches in its wake. When IBM launched its new range of Grid-enabled software and hardware to great expectation in the earlier half of 2002, it announced that Grid computing will take e-business to the next level by giving customers a resilient, flexible, virtual IT infrastructure readily available from any location, on demand.

ALAN FERGUSON, Technical Director at Anix Group questions the usefulness of grid computing for all but the biggest players.



When Sun launched a SuSE Linux version of its *Grid Engine* software in April 2002, it stated that its vision is of a truly heterogeneous grid computing environment in which enterprises can maximise resource utilisation and gain competitive advantage.

Quick returns?

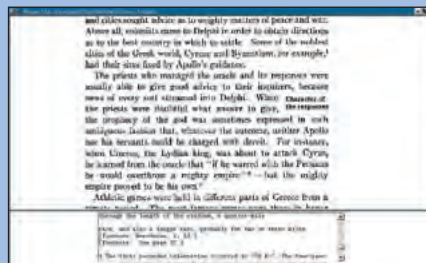
The opportunity to make use of an enormous installed base of processors, with a relatively modest investment and very quick returns, is utterly compelling. Gartner Dataquest reckons that most existing Intel servers only operate at about 5% to 20% of capacity during the workday, and virtually at zero capacity at night, which can only further support the case for grid computing. As processor speeds increase even further, there's going to be even more of this spare capacity to call on in the future.

NON-BUSINESS USES OF GRID COMPUTING

Get involved with these worthwhile projects

PROOFREADING

When we talk about grid applications, it's easy to think of processing power in terms of CPU cycles. But Grid applications can harness other resources too, as the Gutenberg proofreading effort proves. The Gutenberg project aims to deliver written works which have fallen out of copyright into electronic form for ease of access to the general public. Books are usually scanned and subjected to

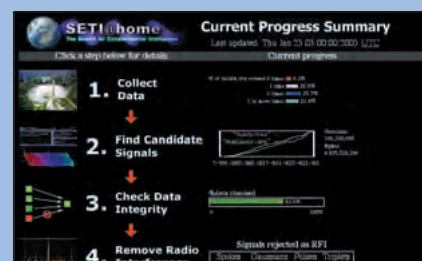


OCR software to translate them to text form. The process isn't foolproof, and the works need to be read through and checked against the scans. This is a grid task that uses the processing power of your brain to read pages on screen and correct the text where necessary. Help the Gutenberg project to safeguard free access to written works by signing up at

<http://charlz.dns2go.com/gutenberg/>

IT PHONE HOME!

For the SETI@Home (Search for Extraterrestrial Intelligence) project, in which PC users worldwide donate unused processor cycles to help the search for signs of extraterrestrial life by analysing signals coming from outer space. We run this in the LXF office; as well as iterations for Windows and Mac, there's a command-line



version for UNIX, WinNT, OS/2, BeOS, Mac OS X Server, OpenVMS, etc to download at <http://setiathome.ssl.berkeley.edu/>

Other medical projects requiring complex molecular analysis have used grid computing, particularly gene mapping, as well as electronic design automation. Grid software also is being increasingly deployed in the financial services industry and corporate research and development departments.

Everyday use

However, despite the claims, and the early adoption in the scientific and technical markets, there is no evidence yet that grid computing is applicable to the needs of common-or-garden IT departments supporting retail, manufacturing, utilities, e-commerce or government organisations – in short, the vast majority of computer users.

A cursory glance at how grids operate should stop dead in their tracks anyone thinking about spending half a million pounds to support their core ERP, financials, sales and marketing applications. First, grids work on the principle of dynamic load sharing, of intelligently allocating parcels of work to available processing capacity across a network. What they don't do is guarantee when that processing task will be finished, let alone the entire submitted project. Guaranteed response times now underwrite the success of all business applications.

Secondly, while allocation of resources is dynamic, the type of data processed on spare CPU cycles is static, making grids suitable for dealing with very large sets of essentially very simple data – in short, for number crunching. And apart from large overnight batch runs, there is no real need for such capabilities in the vast majority of commercial business applications, which are more or less transaction-based. Hence the Grid's appeal for processing terabytes and raw data used in mapping human genomes or modelling air flows in jet engines.

Inevitably, when networks are in place, the issue of security also rears its head, in this case forcing one to ask how safe the underlying file systems are, and addressing the relative insecurity of TCP/IP-networks, on which the Internet is based. And with the Internet the predominant medium for grid computing, bandwidth and availability issues must also be taken into account, issues which don't affect organisations with professionally managed private data networks.

But if Grids are inappropriate for the vast majority of organisations, that doesn't mean the benefits derived from them are unattainable. Making better use of existing resources, for instance, has been the *raison d'être* behind server consolidation, and of the gradual shift towards capacity-on-demand and blade servers.

Ironically, server consolidation centralises the very distributed nature of grid computing to enable many of the same capabilities inside one server or a small cluster. Using spare capacity inside the server also makes use of extremely fast and reliable system buses and storage channels, guaranteeing much higher levels of availability and bandwidth, and all within an environment for which security is geared up to extremely robust specifications.

Server consolidation begets other benefits as well, not least those derived from the process of having to fully audit existing IT infrastructures which over the past five years have expanded with little or no reference to a central plan. The end result is fewer larger and more manageable servers that will be more cost-effective in the long run, even if the up-front server consolidation costs may compare unfavourably with a grid implementation.

And those computers themselves might be made up of hundreds or even thousands of real or virtual servers. One of IBM's propositions, for instance, is the ability to load up

GRID COMPUTING

The newbie need-to-know

WHAT IS GRID COMPUTING?

Grid computing harnesses the power of many (sometimes thousands) of computers to solve problems requiring a large number of processing cycles and involving huge amounts of data that are beyond the processing limits of individual computers. The difference between a grid and a 'cluster' is that cluster tend to be designed to work co-operatively on (usually) a single task with resources controlled centrally. Grid computing works by sharing and aggregating resources distributed across many administrative zones, with all scheduling and resource management handled on the node itself. Grid computing is quite similar to distributed computing, except it has greater scope for

users to finely tune task assignment and co-ordination among the grid elements. Grid computing theoretically allows specific levels of security and access control to be implemented at all levels of the infrastructure, but its dependence on networking across usually open environments such as the Internet can considerably undermine this.

IS SPECIAL HARDWARE OR SOFTWARE NEEDED FOR GRID COMPUTING?

Unlike distributed processing with a central server hub, grid applications typically use a form of peer to peer technology, so all that is required is the grid application, middleware and a network connection to other nodes in the grid.

one of its mainframes with thousands of Linux implementations, each acting as a separate server running a distinct application, and managed by arguably the most proficient set of tools and capabilities available, derived from decades of experience in the data centre.

Grid in a box

In the commercial space, the distributed nature of Grids is further mirrored by the emergence of storage area networks, although like servers the urge to manage them centrally has become that industry's main goal. Either way, the combination of multiple servers in one box and centrally

“There's no need for Grid capabilities in the vast majority of commercial business apps, which are more or less transaction-based”

managed distributed storage clearly point towards one inevitable conclusion: the Grid in a box.

The real grid, as it stands today, represents a niche technology applicable to specialist applications in which availability of raw processing power on demand to attain faster results is far and away the most important implementation criterion. Standards, security and bandwidth, while likely to improve over the next few years, remain secondary considerations to companies and organisations whose key research requirements include crunching inconceivably large amounts of raw data.

And that, in a nutshell, is a million miles away from the transaction-oriented nature of the day-to-day IT needs of most businesses. Grids have a lot to offer a relatively small number of very large organisations. For the rest, the benefits gained by applying the best elements of Grid computing within a structured and centrally managed commercial computing environment are well worth pursuing today. ■

To Protect and Serve

NICK VEITCH introduces the first of our regular sections on storage by looking at one key use of storage – for backing up your data

Whether you have a single desktop or racks of servers, if you have valuable data you need to protect, you need to have a backup strategy. One of the common mistakes of backup policy is to settle on a solution without actually having a strategy. The strategy should define the list of available solutions, not the other way around. Unfortunately, in the world of data backup, the point at which you realise your solution doesn't fit your needs is usually when you lose some critical data, or can't restore your system. You may never have a disaster, but having an ineffective backup solution is as unwise as having an inappropriate insurance policy.

A backup strategy

Before you choose any software, hardware or do anything else, you should start to map out your backup strategy. Your budget and the equipment you have available will certainly have an impact on the strategy you opt for, as is the type of data you are saving and why you are saving it. The first question you should ask yourself is "what do you want from a backup system – what is the ultimate goal?"

If you are saving mission-critical database data, which needs to be restored as quickly as possible in the case of disaster, then your strategy should concentrate on recovery speed and data integrity as primary concerns. If data is important, but not so urgent, you may opt for a different strategy. If you are merely storing data for convenience in case users delete individual files, you might want to opt for a system that is more geared towards individual file recovery than complete system restores. Do you need to backup actual software or any the data created with it? Is it necessary to do complete system backups, or can the data be backed up separately from the OS/software? Is all the vital data secured on fixed systems, or will you need to backup from laptops and other mobile devices?

To properly assess your needs, a good technique is to write out a list of goals: eg how long should it take to restore a system; how much time you wish to devote to backup activity; how long should backups be kept for etc.

It can be convenient, depending on the complexity of your needs, to sketch these out as absolute and desirable values. This should help clarify the priorities your backup strategy should emphasise.

Backing up to what?

Traditionally, backups have been made to removable media formats such as tape drives and even floppy disks. Backing up your fileserver on floppies probably isn't an option these days, but although tape backup systems are still going strong, there are other options available.

Tapes have always been considered the backup method of choice. On the plus side, they offer reasonable value per stored megabyte, and are fairly portable and reusable. There are all sorts of tape options and drives depending on whether speed, cost or reliability is your key concern; it would be quite easy to fill these pages with a discussion of the various tape formats and we probably will in the future, but not today.

The downside for tape backup is that you have to invest in the hardware and the tapes. Tape drives are not cheap, and they aren't particularly versatile – backing up your data is all you are likely to do with them. Tapes themselves, although new technologies have improved reliability, are still very susceptible to a variety of disasters – magnetic, water or flame damage; though paradoxically, if kept well, they probably have a better long-term lifespan than some other solutions. Another minus point is also that they tend to be rather slow, both for backing up and restoring data. 10GB an hour might not sound too bad, unless you have a 1TB fileserver to do daily backups on.

For smaller scale needs, CD-R and DVD-R media might be more suitable. CD-Rs are ubiquitous, cheap, multi-use and quite fast. The downsides are that they are only really suitable for small-scale backups, unless you really like sitting in front of the burner feeding it new discs. DVD-Rs have recently become more inexpensive, with a greater storage capacity, and the media itself is quite cheap. RW media offers the promise of reusability, but of course you have to blank the media first, which is an additional time consideration. Neither format is quite as reliable as their proponents might make out, but they are pretty easy to duplicate too, which might make offsite storage more viable.

Other removable media such as Zip and Jaz disks don't really have the capacity for serious backup tasks.

An emerging solution for backups is to quite simply copy the data onto another hard drive. Hard drives have grown in capacity and dropped in cost to the point where this is probably the fastest growing backup option, even for home users. With 250GB drives costing as little as £250, this can clearly be a cost-effective solution, as well as a convenient one. As well as being fast, hard drives are also random access, which means finding and restoring data is a lot easier too. Even situating the device offsite isn't too problematic. In fact, there are many companies already offering network storage options aimed at the backup market. You could use a hard drive caddy system to physically remove devices to another location, but the drives aren't really robust enough to take this kind of treatment on a daily basis for too long.

“If you're saving mission-critical database data, your strategy should concentrate on recovery speed and data integrity”

Types of backups

To further confuse things, a number of different types of backups are available. Usually backup software will support a number of different schemes, and may even name them differently, but there are three main types.

THE NORMAL OR FULL BACKUP

This is, as you would expect, a complete backup of the chosen filesystem. All the files are saved, and the main advantage of this method is that you only need to restore from one set of backups to get all your data back. This will normally offer the best restore time too, depending on whether you are restoring single files or the whole system. The downside is that these backups take the longest time to perform, and consume the most amount of media.

INCREMENTAL BACKUPS

These seem like a good idea. Instead of backing up all the files on a system, they back up just what has changed since the last backup. This means the backups themselves are fast, and apart from the first one, don't take up as much space. Faster and cheaper are two words that are always in vogue, especially in today's financial climate. However, in most circumstances, this sort of backup probably won't meet your needs. In order to restore a complete set of data, you'll have to restore the initial backup and each incremental backup you have made. If you are merely trying to restore a few selected but critical



The Ecirix VXA1 is a popular tape solution for backups.

ORGANISE FOR EFFICIENCY

Making your admin easier with good system organisation

How you actually store the data and software that needs to be backed up can make a huge difference on your backups, both in terms of complexity and in terms of size. On multi-user systems, one obvious safeguard and convenience is to install the `/home` filesystem in a different partition. This has a number of advantages: If there is a bad filesystem error, then one or other of the OS or the data should be fine; you can easily enforce quotas and you can choose a convenient partition size that makes it easier to backup.

Of course, even without using a different partition you can still separate the tasks of backing up the OS and applications separate from backing up the data. It will depend on your strategy, but

in general the data will be more critical.

For similar reasons it may be useful to partition off server-space data in the `/var` directory, which traditionally includes webserver and database files, as well as logs and other information.

Bear in mind that you don't have to use the same process to back up all the elements of your system. If you can isolate the unchanging aspects of the core OS for example, this could be backed up as a partition image for a quicker restore.

Individual user data may also contain a large number of items which rarely or never change too (documents, images etc) but often the task of separating these out from the volatile data is not worth the effort.

COVER FEATURE **STORAGE****SECURITY**

It's all very well having strict and draconian security policies in the workplace. You might have firewalls, intrusion detection, a rigorous password policy and well-trained staff, but if you don't take care of your backups, all your security measures could be for nothing. Essentially, anyone getting hold of the backups can read any data that you have backed up, no matter how difficult it would be to obtain from the live system.

Backups of sensitive material need to be as closely guarded. If you use an off-site storage option, you are basically putting yourself and your data in the hands of their security procedures.

One crucial thing you can do to ease security worries is to encrypt the data being archived. Most commercial backup tools will give this as an option, but there are plenty of Linux encryption tools even if you are using a home-grown solution.

files, it might not be apparent which backup volume the latest version is on. And even a straightforward system restore could take significantly longer than other backup methods.

If your reason for having backups is merely that you could, if you really had to, recover any file, then this sort of scheme may work. However, for any mission-critical data,

it's not a great deal better than having no backup at all, at least if you allows tens of generations of incremental backups to gather before you do a complete set.

DIFFERENTIAL BACKUPS

A sort of half-way house between a full backup and the incremental. In this type of procedure, only the data changed since the last full backup is stored. This means that you only ever need the full backup and the latest differential backup for a restore, which is safer and faster than using incremental backups. They will take longer and require more space than an incremental backup, but obviously not to the same extent as a full backup.

Testing

Whatever solution you decide on, you really need to test it out, for the same reason that companies have fire drills. If there is a flaw in the process, you want to find it before it is likely to have catastrophic results. Obviously you might not want to restore to a live system, but if you can spare the hardware, setting up a spare and attempting to perform a number of restore tasks could be a useful exercise. ■

BACKUP SOFTWARE

Just some of the software you can use to secure your data

ARKEIA

One of the most popular commercial backup solutions for Linux, Arkeia recently released version 5 of this award-winning software, which we reviewed in issue 37 of *Linux Format*. It handles manual and scheduled backups to a number of devices, including hard disk backup, and a wide variety of tapestreamers and robots. New features include 'Hotbackup' plugins for various database types and a radical new interface. A version for one server and two client machines is available to download free of charge from www.arkeia.com

DANZ RETROSPECT

A big name in data backup, Retrospect now supports Linux clients, although the backup server still needs to run on Windows, which may make it less appealing to Linux-only setups. Retrospect also supports a wide range of devices including tape drives, autoloaders and CD-R/DVD-R burners. See the review in the accompanying issue of *Linux Format* issue 38, and point your Internet browser at www.dantz.com

AMANDA

Stands for Advanced Maryland Automated Network Disk Archiver – a public domain backup utility developed at the University of Maryland. It is as advanced as a free backup

utility gets, and has an active user base.

Amanda works as a single backup server for a number of clients and will work with practically all Linux-supported tape streamers and quite a few autoloaders and robots too. It might not be as simple to configure as many of the commercial options, but does a good job for small to medium-sized backup situations. Visit www.amanda.org

BRU


One of the oldest backup solutions for Linux (the first version appeared for Unix systems before Linux was even born), now owned by the TOLIS group. Has many, many advanced features and is available in a range of versions for different uses from www.tolisgroup.com

NETVAULT

BakBone's flagship backup product has enviable Linux support, with over 20 different distributions supported (a list that is updated regularly and tracks new releases too) and a wide support for Linux supported databases too. A variety of versions to suit most backup requirements up to Enterprise servers and datacentres are available for your delectation at www.bakbone.com

BACULA

A GPL backup solution using the client server model, bacula is a suite of software which can



We featured Netvault in *LXF* issue 14 – since then, there's been many worthwhile updates.

offer an answer to more complicated backup problems. Including such features as autoloader support, pool and volume management, and a disaster recovery mode, this is a fairly comprehensive solution. Again, it does require more knowledge and fiddling to set up, but the results may be worth it.

www.bacula.org


TIME NAVIGATOR

Backup with a *Doctor Who* twist. It allows you to 'step back in time' by managing backup resources and a database for what is stored, you can skip backwards and forwards through time itself. The main advantage of this is that is some critical data becomes corrupt, you can navigate back to the point just prior to when the damage occurred. Supports many devices and network architectures including NAS and SAN. See www.atempo.com

SCSI HARD DRIVE

IBM Ultrastar 146 Z10

IBM's 146.8GB UltraSCSI320 drive offers consistent performance, but is it fast enough? asks **ED HAYNES**.



This is the flagship high capacity drive of IBM's 10,000rpm range, boasting a 146.8GB capacity from a four disk platter mechanism. Like other IBM mechanisms the circuitry on the underside of the mechanism is bare enabling connection to the Ultra320 card via either an 80pin interface such as on this model or the slightly more common 68pin interface connector.

The drive boasts fast seek times averaging 4.9ms and utilises a segmented 8MB buffer to provide data rates of up to 825Mbps across the buffer. The fifth generation head technology enables the drive to read data off the disk with twice the area density of data than the previous

generation at 26.3Gbits/sq. inch. Higher data density means data can be read and written from the platters faster at the same rotational speed.

In operation the drive spins up on the start command and while quieter than some 10K models is fractionally louder than the equivalent Seagate drive at 3.7 Bels. Also the operational volume of the drive is noticeably louder than most IDE mechanisms. Both the drive and the integrated circuits ran warm so sufficient cooling needs to be used to keep the drive well within its operational limits. The construction of the drive is similar to the physical solidness of most IBM hard disks, having a tank-like feel and weighing in at a hefty 780g.

Performance from the drive peaked with read speeds of 66.2MB/sec but averaging at 49.9MB/sec across the entire drive using HD Tach. The write speeds while good peaking at

53.9MB/sec tailed off towards the end of the disk reducing the overall average write performance to 46.5MB/sec.

While the performance is impressive compared to normal IDE drives and many SCSI mechanisms, there are certainly faster drives such as the Seagate 10K.6 featured below. While IBM's 146.8GB drive is marginally cheaper, averaging out at £5.60/GB, the reduced performance is not equal to the lower initial cost of purchase.

INFORMATION

PRICE £823 EX VAT £700
SUPPLIER Dabs.com
PHONE 0870 429 3220
WEB www.ibm.com

SPECIFICATIONS

TYPE UltraSCSI 320
RPM 10,000
SEEK TIME 4.9ms average
CAPACITY 146.8GB
BUFFER 8MB

VERDICT**IBM ULTRASTAR 146 Z10****FOR**

- + Well built and substantial mechanism
- + Good peak read performance

AGAINST

- Write speeds are disappointing
- Slightly noisy during operation

Value 6/10

Features 8/10


Performance 7/10

Overall 7/10

SCSI HARD DRIVE

Seagate Cheetah 10K.6 146.8GB

Seagate's sixth generation Cheetah delivers biting performance without the noise or heat issues that troubles some lesser 10K drives. **ED HAYNES** shows that though it's pricey, you get what you pay for.



Seagate has built a well-deserved reputation with its high performance drives. The Cheetah range shows a mature design and is one of the few which conforms to Ultra320 SCSI specs. Seagate has increased performance with a revised chipset but platters, heads and spindle speed are also fundamental to its 25% higher throughput.

Typically, 10,000 rpm drives suffer from high heat output, not just through high spindle speed but also from the high performance chipset. While the Cheetah certainly doesn't run overly hot, decent cooling should be a consideration. The eight head/four disc

design, makes the drive marginally heavier than many (731g) and the high density platters are more susceptible to damage through knocks when operational. Despite this, the very high platter density means it can be used efficiently in a RAID, requiring physically less drives to create a high capacity server storage solution or for filling network-attached storage boxes.

A considerable bonus of this drive is its low noise level; at just 3.6 Bels it's as quiet as your average 7200 rpm IDE mechanism, which is impressive for any 10,000 rpm drive. Most significant of all is the performance. The Cheetah produced maximum read speeds using HD Tach of 74.4MB/sec and with an average performance across the disk of 53.9MB/sec it's no sloth. The 8MB cache aided the writing, with a max speed of

59.8MB/sec, yet providing great consistency when writing to disk, averaging an impressive 58.4MB/sec, above most of the competition.

With a price of about £5.80/GB this is certainly not cheap compared to the £1/GB of IDE drives, yet it offers unrivalled performance and truly impressive consistency of speed.

VERDICT**SEAGATE CHEETAH 10K.6 146.8GB MODEL ST3146807LW****FOR**

- + High quality drive design and construction
- + Exceptional drive read and write performance

AGAINST

- Read performance reduces at disk's end
- High heat output requires a serious cooling solution

Value 6/10

Features 8/10

Performance 10/10

Overall 8/10

URL Rewriting with Apache

On large sites, particularly when a significant proportion of the content is dynamically generated, URLs often become complex and difficult for users to reproduce. With separate scripts providing different types of content, by virtue of a unique reference number for each item, there becomes a significant amount of repetition in the URLs when users come to type them in. Of course, having shorter and simpler URLs makes the site look much cleaner and professional.

However, we can avoid many of the pitfalls of this by utilising the URL rewriting capabilities of the *Apache* web server. URL rewriting allows us to parse a URL sent to us by a client, and generate an alternative URL, from which the content is sent to the client. With some imagination, it is possible to rewrite a URL into almost anything else we like, without a great deal of effort.

Getting started with URL rewriting

With *Apache*, the most basic form of URL rewriting is with the 'Alias' option, provided by the `mod_alias` DSO. This creates a URL 'symlink', allow us to alias a URL to a particular file or directory. While the **Alias** directive is most commonly used to create links to directories outside of our DocumentRoot, it can also be used to create aliases to any location within our DocumentRoot. Commonly *Apache* has an **Alias** for the `/icons/` folder, so we only have a single folder containing icons, rather than one for each VirtualHost. This **Alias** directive is written as:

```
Alias /icons /usr/local/apache/icons
```

If we then make a request for `/icons/png.gif` on our server,

In the first of a two-part series, DAVID COULSON examines how `mod_rewrite` and virtual hosting with Apache can make life easier for navigating large sites.

it will deliver the contents of the file that is located at `/usr/local/apache/icons/png.gif` on our server. Of course, *Apache* requires a **<Directory>** directive for anything we try to access outside of the main DocumentRoot, otherwise it will use the default options which may, or may not, do what is intended.

Alias only works within a server, so if we want to rewrite the URL and redirect the client to another host, we need to use something a little more special. Much like **Alias**, the **Redirect** directive allows us to redirect to another URL, rather than to a location on our web server. As an example:

```
Redirect /support/ http://support.example.net/
```

Any request made to `/support/index.html` on our server would end up at `http://support.example.net/`, via the use of a HTTP request by *Apache*. However, it will send a 302 response to the client, which is interpreted as a temporary redirect. We can also modify the redirect to send a response to the client saying that our file has been moved permanently, has been replaced by another file, or has been removed from this location. We can either use **RedirectPermanent**, or add a **permanent** flag:

```
Redirect permanent /support/ http://support.example.net/
```

These redirects are fairly basic, so we can only redirect an individual file, or a whole directory. We can't use **Alias** or **Redirect** to change only a specific group of files in a directory, or rewrite the filenames of files within a directory. *Apache* implements regular expression rewriting using **AliasMatch**, so we can easily rewrite a URL exactly as we want. If we want all requests for a GIF file to be redirected to the corresponding JPEG, we can do:

```
AliasMatch /images/(.+).gif$ /home/httpd/htdocs/images/$1.jpg
```

As with the regular **Alias** directive, we still have to provide the complete path to a location on the Unix filesystem.

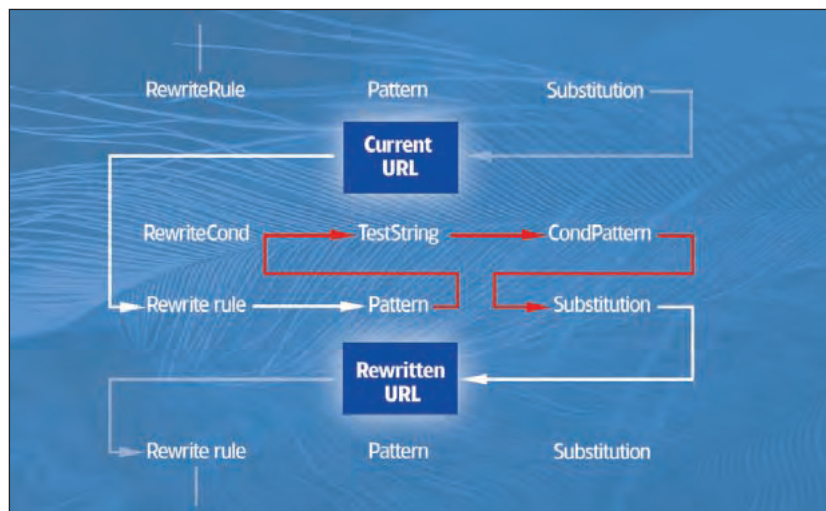
AliasMatch will use a standard regular expression, so almost anything within our URL can be matched and modified before fetching the file from the filesystem.

Likewise, we can also use **RewriteMatch** to rewrite to an external URL via a regular expression:

```
RedirectMatch (.*)gif$ http://www.example.com$1.jpg
```

Alias will only deal with standard files on the filesystem, so if we setup an **Alias** to a CGI, then it will spit back the Perl code, as it isn't sent through `mod_cgi`. As the `/cgi-bin/` directory is generally held outside of the main `htdocs` folder, for exactly this purpose, we need to make use of another `mod_alias` capability in order to rewrite our URL to a CGI. We can use **ScriptAlias** to alias a URL to a location

Figure 1: How the *Apache* URL rewrite process works.



containing a script, although the directory containing our CGIs actually needs to be configured to use the **cgi-script** handler, with:

```
AddHandler cgi-script .cgi
```

Creating the alias is done exactly the same as we would with a regular file:

```
ScriptAlias /cgi-bin/ /home/httpd/cgi-bin/
```

We can also use a regular expression, rather than simply rewriting our directory. If we wanted to rewrite all CGIs to our cgi-bin folder, we can do:

```
ScriptAliasMatch (.*)/(.*)\.cgi$ /home/httpd/cgi-bin/$1/$2.cgi
```

Now, we can access a .cgi file without having to let everyone know they're actually all running within our /cgi-bin/ folder.

While these options are perfectly fine for a couple of virtual hosts, or specific options for a site, if we're running 1000 sites and want them all to handle URLs the same way, we need to find a slightly cleaner way to do it all. *Apache* provides us with **mod_vhost_alias**, which is a virtual host version of **mod_alias**, which is not exactly unclear from the name.

Before using any of the **mod_vhost_alias** options, we need to make each request use the hostname sent to us by the client, rather than that defined in **ServerName**, via the **UseCanonicalName** directive.

```
UseCanonicalName Off
```

We can then setup a virtual **DocumentRoot** directive, which modifies the **DocumentRoot** based upon the hostname used for the request:

```
VirtualDocumentRoot /usr/local/apache/vhosts/%0
```

If we send a request to

http://www.example.net/index.html, our *Apache* server will fetch the file from its location at **/usr/local/apache/vhosts/www.example.net/index.html**.

As always, there is another directive to handle CGIs, so if we want to provide a cgi-bin folder to our users, we can simply do:

```
VirtualScriptAlias /usr/local/apache/vhosts/%0/cgi-bin/
```

Rewriting with mod_rewrite

mod_rewrite is a module for *Apache*, which is distributed as part of the main *Apache* system, which allows for fine grained URL rewriting and redirection. Everything we have looked at so far can be implemented with **mod_rewrite**, but instead of remembering half a dozen different *Apache* directives for **mod_alias** and **mod_vhost_alias**, we can simply use one directive in **mod_rewrite**, **RewriteRule**. **mod_rewrite** handles everything using regular expressions, allowing us to specify exactly what URL components are written, and insert selections from the source URL into our redirected URL.

A basic rewrite with **mod_rewrite** involved enabling the **mod_rewrite** subsystem for our virtual host, then performing the actual rewrite:

```
RewriteEngine On
```

```
RewriteRule /news/([0-9]+)$ /news.php?id=$1
```

This will rewrite **/news/39** to **/news.php?id=39**, by supplying the content of the destination transparently to the client. As far as the client is concerned, it's actually getting

the content from **/news/39**, as it receives no redirect code from our server. If a redirect is needed, such as certain pages have issues being accessed via a different URL by the client, then we can have **mod_rewrite** tell the client that it is redirecting it to a new URL.

```
RewriteRule /news/([0-9]+)$ /news.php?id=$1 [R]
```

The 'R' option, within square brackets refers to **Redirect** and is one of many rewriting options which are available to **mod_rewrite** administrators. Of course, we can do significantly more complex rewrites using **RewriteRule**:

```
RewriteRule ^/~([^\+]+)/?(.*) /u/$1/$2 [R]
```

```
RewriteRule ^/([u|e|g])/([^\+]+)/$ /$1/$2/ [R]
```

This will rewrite any request to '~user' to **/u/user/**, so if we're not using regular public_html directories, we can 'fake' it using **mod_rewrite**. However, it is often useful to check a URL before we perform a rewrite, so we need to

“Apache URL rewriting allows us to parse a URL sent to us by a client, and generate an alternative URL if needed”

take advantage of the **RewriteCond** directive. This will check one value against another, such as **%{REQUEST_FILENAME}**, or any other environmental variable, against a regular expression. If this is true, then the following **RewriteRule**, or indeed another **RewriteCond**, will be performed by *Apache*, and if it false, the **RewriteRule** after the 'true' **RewriteRule** will be used.

To demonstrate this, **mod_rewrite** can point your browser at the default hostname for it:

```
RewriteCond %{HTTP_HOST} !^fully.qualified\.domain\.  
.name [NC]
```

```
RewriteCond %{HTTP_HOST} !^$
```

```
RewriteCond %{SERVER_PORT} !^80$
```

```
RewriteRule ^/(.*) http://fully.qualified.domain.name:  
%{SERVER_PORT}/$1 [L,R]
```

```
RewriteCond %{HTTP_HOST} !^fully.qualified\.domain\.  
.name [NC]
```

```
RewriteCond %{HTTP_HOST} !^$
```

```
RewriteRule ^/(.*) http://fully.qualified.domain.name/  
$1 [L,R]
```

RewriteCond first checks the **HTTP_HOST** variable, which is the value of the **Host** header sent by the client. If it is already that of our server, then it simply ignores it.

RewriteCond then checks to ensure that **HTTP_HOST** is actually set and to find out if the port used is something other than 80. If all of these conditions are true, then it will proceed to rebuild the URL around our domain name and port, and redirect the client to it. The **L** option will tell **mod_rewrite** that this is the last rewrite rule to be used, assuming it checks out to be 'true'. The next section performs exactly the same function, except it does not add on the **SERVER_PORT** variable. The **[NC]** option ensures that the check against **HTTP_HOST** is not case sensitive. This should not be confused with the 'next' and 'chain' options, which would be written as **[N,C]**, rather than **[NC]**.

RewriteCond can also do many other interesting checks, so we can have it check to see



if the file exists or do a 'look ahead' on the URL to see if it is valid.

```
RewriteEngine on
RewriteCond /your/docroot/%{REQUEST_FILENAME} !-f
RewriteRule ^(.+) http://webserverB.dom/$1
```

```
RewriteEngine on
RewriteCond %{REQUEST_URI} !-U
RewriteRule ^(.+) http://webserverB.dom/$1
```

The first example checks if a file exists, and if it does not, modifies the URL to that of another web server transparently. The latter will check to see if

`%{REQUEST_URI}` will return a file via the server, and if it does not, then it will perform the redirect. Both of these are the equivalent of a 404 document, although the example which uses `%{REQUEST_URI}` requires an internal request, which needs a little CPU time.

Similarly, we can have our server search multiple DocumentRoots for our requested document, effectively creating search 'layers'.

```
RewriteEngine on
# first try to find it in custom/...
# ...and if found stop and be happy:
RewriteCond /your/docroot/dir1/%{REQUEST_FILENAME} -f
RewriteRule ^(.+) /your/docroot/dir1/$1 [L]
# second try to find it in pub/...
# ...and if found stop and be happy:
RewriteCond /your/docroot/dir2/%{REQUEST_FILENAME} -f
RewriteRule ^(.+) /your/docroot/dir2/$1 [L]
# else go on for other Alias or ScriptAlias directives,
# etc.
RewriteRule ^(.+) - [PT]
```

A more useful example of `mod_rewrite` is to redirect ~user requests which are made by users not within our

“mod_rewrite has some very powerful capabilities to produce a dynamic virtual hosting platform without using Perl”

DNS domain, and send them off to another location. There are many reasons which they are useful, particularly within an office or educational establishment where internal and external versions of pages should be maintained by users.

```
RewriteEngine on
RewriteCond %{REMOTE_HOST} !^.\ourdomain\.com$
RewriteRule ^(/~.+) http://www.somewhere.com/$1 [R,L]
```

So far we've only managed rewriting static content with `mod_rewrite`, as if we rewrite to a CGI, it won't be handled by the `cgi-handler` and `mod_cgi`. Instead, we need to use `mod_rewrite` to tell *Apache* that this script needs to be sent through `cgi-handler`, which is done by modifying the MIME

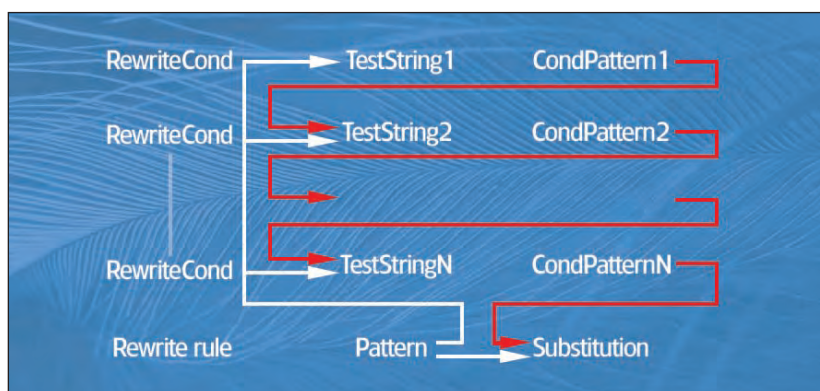


Figure 2: the Rewrite substitution process

type of the document. As CGIs sent their own MIME type as part of the headers, modifying the MIME type has no effect on the processing of the CGI and is transparent to the client, as well as the individual who wrote the URL.

```
RewriteEngine on
RewriteBase /~quux/
RewriteRule ^(.+)\.html$ $1.cgi [T=application/x-httpd-cgi]
```

Anything which has its MIME type set to `application/x-httpd-cgi` will be handled by the CGI subsystem within *Apache*, although we do need to permit `ExecCGI` on that location. We can supply any extension we desire, so we don't need to use an extension for our CGI scripts which is consistent with the `AddHandler` directive for `cgi-handler`.

The use of CGIs with `mod_rewrite` can be extended further by creating a CGI which writes its output to a `.htm` file, as well as to `STDOUT`. While this is not ideal for dynamic content, it is very useful for a CPU-hungry CGI which can have its results cached

```
RewriteCond %{REQUEST_FILENAME} !-s
RewriteRule ^page\.html$ page.cgi
[T=application/x-httpd-cgi,L]
```

We could quite easily create a CGI which checks the timestamp on `page.html`, then removes or overwrites it if `page.html` is older than a set period of time. On a similar note, we can also inspect a couple of environmental variables containing the current time in hours and minutes, which we can use to decide what content to send on to the client:

```
RewriteEngine on
RewriteCond %{TIME_HOUR}%{TIME_MIN} >0700
RewriteCond %{TIME_HOUR}%{TIME_MIN} <1900
RewriteRule ^foo\.html$ foo.day.html
RewriteRule ^foo\.html$ foo.night.html
```

Any request that is made between 7am and 7pm will result in `foo.day.html` being sent back to the client, and at any other time of the day, it will return `foo.night.html`. This can be particularly useful to supply out-of-hours contact information, or ensure that clients or users know that they're not going to get someone on the end of a phone at 3am.

Virtual hosting with mod_rewrite

We've looked at some rewriting capabilities of `mod_rewrite`, but it's still not useful to us as a replacement for `mod_vhost_alias`. However, `mod_rewrite` has some very powerful capabilities allowing us to produce a dynamic

virtual hosting platform without having to learn Perl or play with databases. This is achieved by using the **RewriteMap** directive, which makes use of a very simple key/value array, generally supplied by a plain text file, which we can use for lookups and process the results.

A simple use of **RewriteMap** is to rewrite URLs to specific users and groups.

```
RewriteEngine on

RewriteMap user-to-host txt:/path/to/map.user-to-host
RewriteMap group-to-host txt:/path/to/map.group-to-host
RewriteMap entity-to-host txt:/path/to/map.entity-to-host

RewriteRule ^/u/([^\+])?(.*) http://${user-to-host:$1|server0}/u/$1/$2
RewriteRule ^/g/([^\+])?(.*) http://${group-to-host:$1|server0}/g/$1/$2
RewriteRule ^/e/([^\+])?(.*) http://${entity-to-host:$1|server0}/e/$1/$2

RewriteRule ^/([u]ge)/([^\+])?$ /$1/$2/www/
RewriteRule ^/([u]ge)/([^\+])/([^\+].+)$ /$1/$2/www/$3\
```

This would lookup the second directory from the URL in one of three maps, depending upon the first directory used. If we hit **/u/david/test**, it would lookup the value of the key **david** in the user-to-host table. Of course, we can supply a fallback value, in this case **server0**, which is used if there is no corresponding key within our map.

Adapt maps

So far this is all pretty basic stuff, and unfortunately isn't all that much use for virtual hosting. We need to adapt the maps to handle domain names, rather than directories, so we can associate a particular hostname with a user. This can all be done with a single map, which contains hostname and their corresponding DocumentRoot directories.

Our map would look something similar to the following, assuming that *Apache* is running on port 80.

```
##
## vhost.map
##
www.vhost1.dom:80 /path/to/docroot/vhost1
www.vhost2.dom:80 /path/to/docroot/vhost2
:
www.vhostN.dom:80 /path/to/docroot/vhostN
```

We can then modify our *httpd.conf* file, ensuring that all requests to any virtual host will hit our code. We also need to modify the logging somewhat, so that our *access_log* file will successfully log which virtual host each request was for.

```
##
## httpd.conf
##
:
# use the canonical hostname on redirects, etc.
UseCanonicalName on
```

```
:
# add the virtual host in front of the CLF-format
CustomLog /path/to/access_log "%{VHOST}e %h %l %u %t \"%r\" %>s %b"
:

# enable the rewriting engine in the main server
RewriteEngine on

# define two maps: one for fixing the URL and one
# which defines
# the available virtual hosts with their corresponding
# DocumentRoot.
RewriteMap lowercase int:tolower
RewriteMap vhost txt:/path/to/vhost.map

# Now do the actual virtual host mapping
# via a huge and complicated single rule:
#
# 1. make sure we don't map for common locations
RewriteCond %{REQUEST_URI} !/commonurl1/*
RewriteCond %{REQUEST_URI} !/commonurl2/*
:
RewriteCond %{REQUEST_URI} !/commonurlN/*
#
# 2. make sure we have a Host header, because
# currently our approach only supports
# virtual hosting through this header
RewriteCond %{HTTP_HOST} !^$
#
# 3. lowercase the hostname
RewriteCond ${lowercase:%{HTTP_HOST}}[NONE] ^(.+)$
#
# 4. lookup this hostname in vhost.map and
# remember it only when it is a path
# (and not "NONE" from above)
RewriteCond ${vhost:%1} ^(/.*)$
#
# 5. finally we can map the URL to its docroot location
# and remember the virtual host for logging purposes
RewriteRule ^(/.*)$ %1/$1 [E=VHOST:${lowercase:%{HTTP_HOST}}]
```

What would normally happen is that the logs would be piped through *cronolog*, so we end up with a log file per day. A simple script can be used to parse this out and use our map to figure out which virtual host the logs is associated with, so we can write our user their own logs file each night.

Going further

mod_rewrite is a very complex beast, and contains a wide variety of options, all of which can be customised to the administrators specific requirements.

<http://httpd.apache.org/docs/misc/rewriteguide.html> is where the author of *mod_rewrite* has written a particularly useful guide to his *Apache* module.

Next month, we're going to continue looking at *mod_rewrite* and integrating it with scripts and databases to create a far more powerful virtual hosting system. ■

NEXT MONTH

In the next issue of *Linux Pro*, we're going to expand on the database aspect, and we will also be employing *mod_perl* as a replacement for *mod_rewrite* to create completely dynamic hosting systems through the use of embedded Perl code in our *Apache* configuration.

Deploying Java Servlets with Tomcat



For those of you who have never heard of it, *Tomcat* is a specialised kind of web server, written in Java. I suppose we would think of the primary role of an “ordinary” web server as serving up HTML files and GIF or JPEG images.

Tomcat can deliver that kind of static content, but its real purpose is to act as a container for “servlets” written in Java. These servlets make up the server side of a distributed application. Because the communication between the clients and servers in these distributed apps sits atop the HTTP protocol, such apps are sometimes called “web applications” and *Tomcat* might legitimately be called a “kind of web server”.

You’ll sometimes hear *Tomcat* referred to as *Jakarta-Tomcat* or even *Apache-Tomcat*. “Jakarta” refers to all the Java-related projects hosted by the Apache Software Foundation, of which *Tomcat* is only one. (Others include *Ant*, a build tool which supersedes *make* in the Java world, *Lucene*, a search engine, and *James*, a mail server.) Even more confusing, the current implementation of *Tomcat* is actually called *Catalina*.

Java servlets can readily talk to “enterprise services” such as databases (via *JDBC*) or email systems, (via *JavaMail*) or *EnterpriseJava Beans*, or other distributed objects (via *RMI* or *CORBA*), and so the emphasis is mainly on using the servlet technology as the “middle tier” of a classical three-tier distributed architecture, as shown in **figure 1**.

A couple of asides are in order here. First, I would love to know what the term “enterprise services” actually means, since it leaves me unenlightened. Second, I sometimes encounter people who believe that “three-tier architectures” are the only way to build distributed systems. I find this narrowness of view a bit irritating, having spent a good part of my research career building distributed architectures which are decidedly not “three-tiered”. Nonetheless, three-

PART 1 Outside it is damp, cold and dismal, but inside your computer is a world of activity and light, order and great intricacy. Join **DR. CHRIS BROWN** in the warm to take a look at *Tomcat*.

tier is the way most people are building commercial, distributed applications these days, and it is certainly where most of the action is with servlets.

Compiling on the fly

You’ll notice from **figure 1** that *Tomcat* can also act as a container for Java Server Pages (JSP). These are pages of HTML with chunks of java code embedded in them inside special tags, in much the same way that you might embed PHP code into HTML templates using *mod_php* with *Apache*. *Tomcat* compiles these chunks of Java on the fly using a compiler called *Jasper*. In contrast, servlets are pure Java code, rather more like CGI scripts, although this comparison should not be pushed too far, because whereas a CGI script runs in a separate process and is left pretty much to its own devices, servlets run in a context provided by *Tomcat*’s servlet container, and receive considerable support from it.

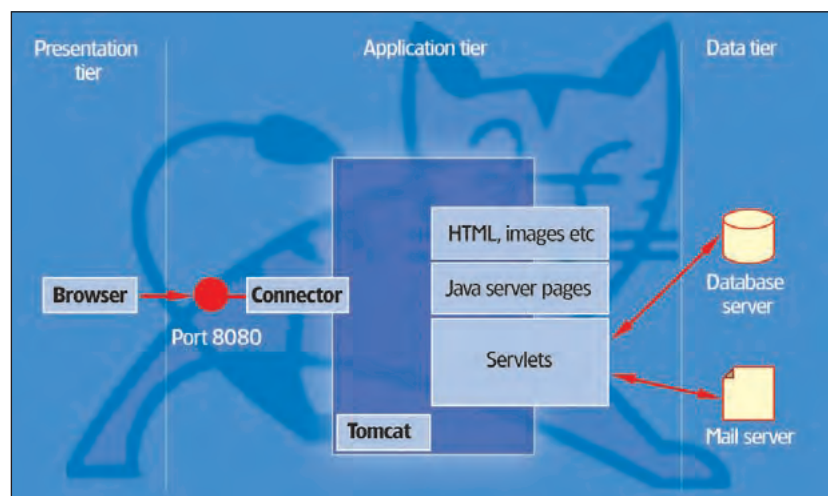
Tomcat is a “reference implementation” – its primary design goal is precise conformance to the standards established by Sun for the environment in which JSP and servlet code executes. Specifically, *Tomcat* v4 provides conformance to v1.2 of the JSP specification and v2.3 of the servlet API specification. (These are available from <http://java.sun.com/products/jsp> and <http://java.sun.com/products/servlet> respectively.) *Tomcat* is by no means the only servlet container available; commercial products include BEA *Weblogic* and IBM *WebSphere*. In the Linux world however, *Tomcat* has the distinct advantage of being free.

Prior to the availability of *Tomcat*, to provide support for Java servlets in the Linux / *Apache* world, you would have used *JServ*, which was a servlet container implemented as a module that integrated directly into *Apache*. *Tomcat* effectively renders *JServ* obsolete. *Tomcat* is not an *Apache* module, it is a stand-alone product capable of functioning as a full-fledged web server in its own right, using a component which *Tomcat* calls an “HTTP connector”. However, you can “hide” the *Tomcat* server behind an *apache*-hosted website, and have *apache* automatically forward parts of the URL namespace to *Tomcat*. In the remainder of this month’s tutorial, we look at what it takes to get a basic *Tomcat* server installed and how to deploy a simple web application.

Installation

Usually, my preference is to install from source tarballs, but the instructions for building *Tomcat* from source run to a total of 27 steps and look a little challenging, to say the least. It turns out that a full build of *Tomcat* requires the build tool *Ant* (also part of the Jakarta project) and more

Fig 1: *Tomcat* as part of a 3-tier architecture.



than a dozen other things besides. In light of this, you may prefer to install from a pre-built binary distribution from <http://jakarta.apache.org> by following the links to "Binaries" then "Tomcat4.1.18" (the current version at the time of writing.) Download the file *tomcat-4.1.18-LE-jdk14.tar.gz* and park it somewhere that you find convenient (I put mine into /tmp).

Installation is then a simple matter of unpacking the archive into whatever directory local conventions suggest is appropriate. I put it into /opt, like this:

```
# cd /opt
# tar zxvf /tmp/tomcat-4.1.18-LE-jdk14.tar.gz
```

You'll discover that you now have a directory called /opt/jakarta-tomcat-4.1.18-LE-jdk14 which contains the installation. To make this a little easier to deal with, I put in a symbolic link called "tomcat", like this:

```
# ln -s jakarta-tomcat-4.1.18-LE-jdk14 tomcat
```

You need the JDK, too

In order to run the server, and to build and deploy the Java web applications that it exists to support, you'll also need the Java Development Kit (JDK) installed. You may have this already, or it may be on your Linux distro's CDs. If not, you can get it from <http://java.sun.com/j2se> by following the links to "J2SE Downloads" and fetching the self-extracting archive file *j2sdk-1_4_1-linux-i586.bin* (beware – this is a 42 MB download!). Put it into /tmp then install it by running it as a shell script, like this:

```
cd /opt
sh /tmp/j2sdk-1_4_1-linux-i586.bin
```

This will create the J2SDK installation in the directory /opt/j2sdk1.4.1

Finally, you'll need to set a couple of environment variables by adding entries to (for example) your .bash_profile file, something like this:

```
export JAVA_HOME=/opt/j2sdk1.4.1
export CATALINA_HOME=/opt/tomcat
PATH=$PATH:$JAVA_HOME/bin
```

Taking it out for a spin

Time to try it out. The *Tomcat* distribution comes with its own special startup scripts and the server can easily be started manually like this:

```
cd /opt/tomcat/bin
./startup.sh
```

If you look at the processes running in the machine (use **ps**) you should now see a whole bunch of Java processes which are the Java virtual machines running the *Tomcat* server. If you look at the active TCP/IP endpoints (use **netstat -ant**) you should see that there is a TCP endpoint in the "LISTEN" state on port 8080 (the default port for *Tomcat*). If all appears well, start a browser and enter the URL **http://localhost:8080**. Hopefully, you will be rewarded with the *Tomcat* home page. Note that *Tomcat* listens on port 8080 by default.

Writing a web application

Once you have *Tomcat* up and running, chances are you'll get sidetracked for quite a while by browsing the documentation or following the links from the home page

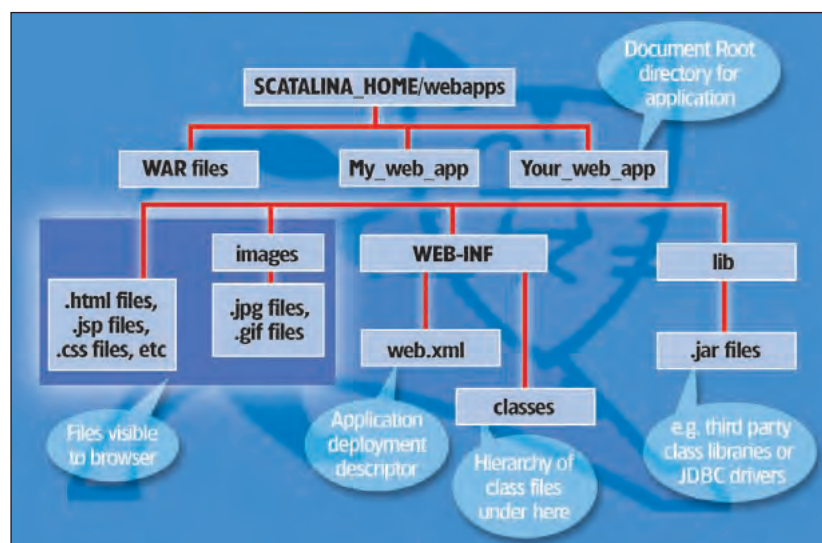


Fig 2: Web Application Deployment Hierarchy.

JAVA BYTE CODE

Processor-independent distribution

With Java, you don't compile programs down to machine code for a real processor (such as a Pentium or a SPARC), you compile them down to something called **Java Byte Code**, which is machine code for a (hypothetical) machine called a Java Virtual Machine (JVM). The Java Byte Code is placed into files called .class files. At run time, a JVM is started to interpret the Java Byte Codes. On Linux, the current version of Java (1.4) uses Sun's *HotSpot* virtual machine which does "Just-In-Time" (JIT) compilation of the byte codes into real machine code for the underlying processor. This speeds execution up considerably. However, the important point is that when Java applications are deployed, it's the Java Byte Codes (ie the .class files) which are distributed. This means that the distribution is independent of the underlying processor.

to try out the sample JSP pages and servlets which are included in the distribution. (The documentation is excellent, by the way.) Sooner or later, however, you may even get the urge to write and deploy a Java web application of your own.

The web application we'll use in this tutorial is about as simple as you can get. It consists of two files. The first is a simple HTML file that defines a form with a text input field called "name" and a button labelled "Submit". When the form is submitted, the form data is sent to our second file which is a Java servlet. The servlet simply echoes the form data back to the browser. Not terribly exciting, but think of it as the "hello world" of java servlets.

Here's the HTML file, which we'll call **formexample.html**. Note that the line numbers in these examples are just for reference and clarity, they are not part of the file.

```
1 <html>
2 <head>
3 <title> Simple Form Example </title>
4 </head>
5
6 <body>
7 <h1> Simple Form Example </h1>
8 <form action="/formexample/servlet/FormExample"
  method="GET" >
9   Enter your name:
10  <input type="text" name="name" />
11  <input type="submit" name="Submit"
    value="Submit name" />
12 </form>
13 </body>
14 </html>
```

Since this isn't an HTML tutorial, I won't attempt to explain this file in any detail. However, you should notice the action attribute of the **<form>** tag at line 8. This URL (/formexample/servlet/FormExample) identifies the Java servlet to which the form is submitted. We'll see how this association is made later.

The source code for our servlet, which we'll call **FormExample.java**, is shown below and overpage:

```
1 import java.io.*;
2 import javax.servlet.*;
```

```

3 import javax.servlet.http.*;
4
5 /* Very simple servlet to access fields of a form */
6
7 public class FormExample extends HttpServlet {
8     public void doGet(HttpServletRequest request,
9         HttpServletResponse response)
10        throws IOException, ServletException
11    {
12        response.setContentType("text/html");
13        PrintWriter out = response.getWriter();
14
15        out.println("<html> <head>");
16
17        String title = "This is the FormExample servlet";
18
19        out.println("<title>" + title + "</title>");
20        out.println("</head> <body>");
21        out.println("<h1> Hello " +
22            request.getParameter("name") + "</h1>");
23        out.println("</body> </html>");
24    }
25 }

```

A little explanation is probably in order. This description is going to make more sense if you know a bit about Java, but even if you don't, you should get the general idea.

The easiest way to write a servlet is to write a class which extends the base class `javax.servlet.http.HttpServlet`. This base class provides a method which retrieves HTTP requests. It reads the HTTP method from the request header and dispatches the request to an appropriate handler. For an HTTP GET request, the handler is called `doGet()`, for a POST request, it's called `doPost()`, and so on. All you have to do is write the handlers for the methods you'll be using. In our case, the form data is passed back using the GET method (see line 8 of `formexample.html`) so we need to provide a `doGet()` handler.

Details, details...

Our derived class, `FormExample`, is defined at line 7. The name of the class must agree with the name of the java file that defines it. Lines 8-10 are the header of our `doGet()` method. When an HTTP GET request is received for a URL which maps onto this servlet, this function will be called. It is passed two objects, of type `HttpServletRequest` and `HttpServletResponse`.

`HttpServletRequest` encapsulates the HTTP request message and provides convenient methods for accessing request headers and form data. You'll see, for example, at line 22 a call to `request.getParameter("name")` which retrieves the value of the `name` field in the submitted form. (It doesn't get much easier than that!)

`HttpServletResponse` provides methods for constructing the HTTP response message, including methods for setting values for the HTTP response headers, and the MIME type of the content. (The latter is essential; see line 12.) It also contains an output stream for sending data back to the client. This is retrieved via the `getWriter()` call at line 13.

The bulk of the `doGet()` function (lines 15 – 23) merely generates some simple HTML output which is returned to

TOMCAT TIP

bash filename completion

If you're typing these hideously long file names out in full, you need to spend a couple of minutes learning how to use the filename completion mechanism in `bash` – over time, it will save you hours.

the client, by calling the `println()` method of the `PrintWriter` object obtained from `HttpServletRequest`. In a more worldly servlet, we would see here some business logic being implemented, and probably access to a database or some remote objects.

Deploying the web application

So much for the code. Now it's time to deploy our application into the *Tomcat* servlet container. Applications can be deployed in two ways. First, they can be deployed "unpacked" by placing the files into a specific directory hierarchy. This is convenient during development, and is the method we'll use. Alternatively, applications can be packaged up into `.WAR` files. We'll get back to that later.

The first thing to do is to compile our servlet into a `.class` file. For this we use the Java compiler, `javac`:

```
javac -classpath /opt/tomcat/common/lib/servlet.jar
FormExample.java
```

Assuming there are no compilation errors, this will create the file `FormExample.class`.

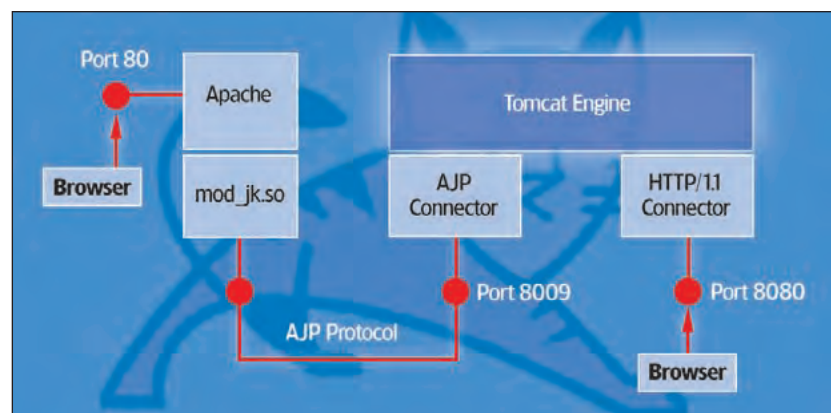
Figure 2 shows the file hierarchy for deployment of a Java web application. The hierarchy shown under the directory `formexample` shows the standard layout defined by Sun's servlet specification. The diagram shows a more complete hierarchy than is required for our specific application, which has only two files: (1) The HTML form, which we'll place in `formexample/formexample.html`, and (2) the `.class` file for the servlet, which we'll place in `formexample/WEB-INF/classes/FormExample.class`. The basic idea of this hierarchy is that files that are intended to be directly visible to the browser are placed directly within the application's top level directory (`formexample` in this instance), whereas the servlets and other supporting resources that describe the application are placed under `WEB-INF`.

There is one other file that, according to the servlet specification, is required to be present. It's known as the web application deployment descriptor; its name is `web.xml`, and it should be placed in the `WEB-INF` directory. This file describes the servlets and other components that make up the application. In a simple case, you can actually manage without this file.

Automatic deployment

The top level directory of our application is placed under the directory `$CATALINA_HOME/webapps`. (If you've installed *Tomcat* into `/opt/tomcat`, as suggested earlier in

Fig 3: Connecting Apache to Tomcat.



the tutorial, this will be the directory `/opt/tomcat/webapps`.) The `webapps` directory is special, in the sense that *Tomcat* will automatically deploy any applications it finds in that directory when it starts up. It's possible to deploy applications into other directories, but then you need to define a "context" for the application, which requires changes to *Tomcat*'s configuration file – something we're trying to avoid for now.

OK, so let's create the required directory structure and deploy our application. The commands below assume that you're currently in the directory where you created your `formexample.html` and `FormExample.class` files.

```
mkdir -p /opt/tomcat/webapps/formexample/WEB-INF/classes
```

```
cp formexample.html /opt/tomcat/webapps/formexample
```

```
cp FormExample.class /opt/tomcat/webapps/formexample/WEB-INF/classes
```

It's time for a short digression to tell you about an interesting hoop I had to jump through to get this example to work on the current version of *Tomcat* (4.1.18). It had previously worked fine on version 4.0.3, but flatly refused to find the `FormExample` servlet under the new version. To cut a long story short (and completely miss out the bit where I went into the garden and pulled the ears off a squirrel in frustration), because I'm not using a `web.xml` file for the application, I am relying on the services of a special servlet called the *invoker*, which is responsible for starting up "anonymous" servlets. The *invoker* servlet is configured in the file `/opt/tomcat/conf/web.xml`. In particular, this file contains the following lines:

```
<servlet-mapping>
  <servlet-name>invoker</servlet-name>
  <url-pattern>/servlet/*</url-pattern>
</servlet-mapping>
```

which establish a mapping such that URLs of the form `http://<hostname>/servlet/*` are referred to the *invoker* servlet for processing. However, in the out-of-the-box *Tomcat 4.1.18*, these four lines are commented out. If you're going to try out the example in this tutorial, you'll need to remove the comment delimiters.

The moment of truth

You'll need to stop and restart the *Tomcat* server to get your application deployed. (Later, you will discover there is an application called *the manager*, which lets you start and stop applications without restarting *Tomcat* each time.)

```
cd /opt/tomcat/bin
```

```
./shutdown.sh
```

```
./startup.sh
```

Finally, in your browser enter the URL

```
http://localhost:8080/formexample/formexample.html
```

Enter a name into the field on the form, click on the submit button, and verify that the servlet is called.

Deploying a WAR file

Deploying a web application as an unpacked hierarchy of files, as we just did, is probably the most convenient approach during development, since it allows individual components of the application to be easily replaced. If a completed application is to be copied onto production

FURTHER INFORMATION

Specification and further reading

If you want to read the full Java Servlet and JSP Specifications from Sun, see <http://www.jcp.org/aboutJava/communityprocess/final/jsr053>. Sun require you to accept a licence before download.

If you'd like a good book about *Tomcat*, we recommend *Professional Apache Tomcat* by Vivek Chopra, Sing Li, Ben Galbraith, published by Wrox Press, ISBN 1-8610-0773-6.

Intended for professionals working with Java web applications, it assumes a working knowledge of the JSP and Servlet technologies, but only from an administrator's point of view. Familiarity with databases, XML, HTML, networking, and general admin techniques is also assumed. There are some sample chapters available to download at www.wrox.com/books/1861007736.htm

servers for deployment, however, it's more convenient to use a WAR file. A WAR file is simply an archive containing the files which make up the application, organised in the standard directory hierarchy we've just examined. It's built using the Java archiving tool, *jar*. (If you're familiar with the linux utility *tar*, then *jar* works much the same way. However, whereas *tar* is used as a general-purpose archiver, *jar* is mostly used to package Java applications, applets and servlets for ease of deployment.)

Once the WAR file is created, it is deployed simply by dropping it into the appropriate directory. You don't even need to unpack it – *Tomcat* will do that, and deploy the application for you. Once you have the application deployed and debugged in its unpacked form, creating the WAR file is easy. In our example, we could do the following:

```
cd /opt/tomcat/webapps/formexample
```

```
jar cvf /tmp/formexample.war *
```

If you want to try deploying the WAR file, you'll need to delete the unpacked version of the application first:

```
rm -rf /opt/tomcat/webapps/formexample
```

Now copy the WAR file into place and restart the server:

```
cp /tmp/formexample.war /opt/tomcat/webapps
```

```
/opt/tomcat/bin/shutdown.sh
```

```
/opt/tomcat/bin/startup.sh
```

You should find that *tomcat* will automatically deploy the application from the WAR file and recreate its directory structure under `webapps`.

Integrating Tomcat with Apache

Although as we've mentioned earlier in this article, *Tomcat* is capable of functioning as a stand-alone web server, it's also possible to link it with *Apache*, so that your entire web site is fronted by *Apache*, which is configured to forward requests for specified parts of the URL namespace to *Tomcat* for processing.

On the *Tomcat* end of things, you use of a component called an AJP connector, and on the *Apache* end, you need to load the module `mod_jk.so`. **Figure 3** shows the overall arrangement. In the example shown, *Tomcat* also continues to listen for HTTP requests on port 8080. AJP stands for Apache JServ Protocol; this is the protocol which *Apache* and *Tomcat* use to communicate. We'll take a more detailed look at the *Tomcat* end of this next month when we get to grips with *Tomcat*'s XML-based configuration file. ■

NEXT MONTH

In the second part of this introduction to *Tomcat* in April 2003's *Linux Pro*, free with issue 39 of *Linux Format*, we'll investigate the mysteries of *Tomcat*'s configuration files.

Network security iptables

PART 1 In the first of a two part series on iptables, **DAVID COULSON** looks at installation and configuration.

Network security has always been a major concern for businesses with online services. With DSL, cable and other 'always-on' technologies becoming the norm, home and small businesses are becoming more vulnerable to attacks and the possibility of having a system exploited. In a perfect world, we could leave unprotected systems online all of the time without feeling threatened by outsiders, but unfortunately the Internet is far from perfect, and there are individuals out there who are happily scanning ISP net blocks for exploitable systems in search of their latest target.

While it's important to keep software up-to-date and limit the amount of potentially exploitable software from our systems, it is also equally important to limit what is available to the outside world. Of course, even the best administrators occasionally make mistakes and unfortunately, there are occasions when a system may be more open than is desired due to a configuration fault.

A firewall can be used to filter out packets heading from the outside world to our network, which would normally allow someone to connect to a network service running on our system. It's worth remembering that, generally, a firewall is a layer 3 and 4 device, so it operates at the IP and TCP/UDP level. This means that a firewall is unable to filter based upon contents of packets. For example, a firewall does not filter email for viruses or check for specific unwanted content within web documents. Usually something more specific to a firewall is required for these capabilities, although the actual software used may be run on the same system as the firewall. Notwithstanding this, a firewall is the first line of defense against an outside attack, so it's well worth taking some time to set it up correct and ensure that our system is as secure as it can be.

Firewalling with Linux

Linux has had firewalling capabilities for a very long time. Back in the 2.0 kernels, we had the *ipfwadm* package, which did basic firewalling. With 2.2, *ipchains* was implemented and offered a significant number of improvements over 2.0. Now, with the 2.4 kernel, all of the old firewalling code was ripped out and a new system implemented. Rather than simply provide firewalling

capabilities, they decided to simply write a system which can take plug-ins, allowing any firewalling system to be used with Linux. This subsystem is known as *netfilter*, and allows *ipfwadm* and *ipchains* to be used with the 2.4 kernel series using the ports of each of these to *netfilter*. However, a far more powerful firewalling system was implemented, known as *iptables*, which not only has far more capabilities than earlier efforts, but it can also handle having these extended through the use of kernel modules, making for a very flexible installation.

We're going to start off by assuming that the Linux system which will be our firewall already has a 2.4 kernel running happily and that building a kernel is not too much of an issue. All distributions which use 2.4 should have *iptables* available, so unless we're using our own kernel, or want to add some extra capabilities to *iptables*, then we won't need to rebuild the kernel.

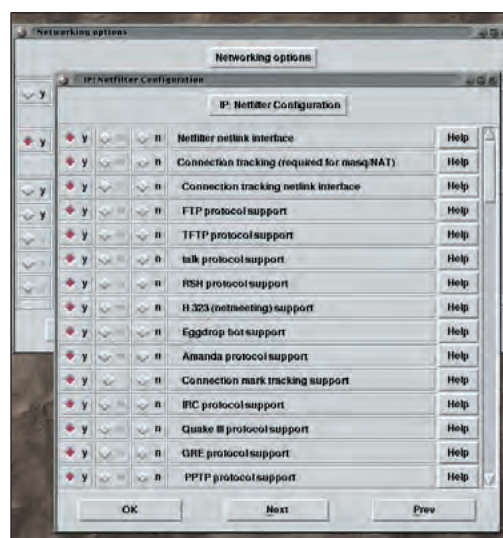
All of the *netfilter* options live within 'Network options', under IP: Netfilter Configuration. Generally, we'd pick all of the options and compile them into the kernel, except for *ipchains* and *ipfwadm*. Having *iptables* capabilities available as loadable modules can make life a little more complex, plus it is rare to actually have the system up and running without the firewall enabled, so nothing is really being achieved by having them set up as modules. However, some people do like having modules, and *iptables* will automatically load the module as needed when a rule is installed which uses that particular module. The only other reason to compile everything for modular use is if we already have a 2.2 or 2.0 firewall setup, and want to continue with *ipchains* or *ipfwadm* under 2.4 before upgrading to *iptables*. It's not possible to compile both *iptables* and *ipchains* into the kernel, for obvious reasons, but we can build them both as modules and 'modprobe *ipchains*' if we want to use *ipchains*.

Once we've picked what we need to compile into the kernel, we simply need to run the standard **make bzlib0** and wait for the kernel to install. If needed, **make modules** && **make modules_install** should be performed in order to install the new *iptables* modules which were compiled. After a reboot, a quick look at the boot logs with **dmesg** should show:

ip_tables: (C) 2000-2002 Netfilter core team

Of course, if it's available as a module, then one needs to **modprobe ip_tables** before checking for that boot message. If for any reason that message doesn't appear, then it's worth checking the kernel configuration and that

Netfilter and iptables have lots of configuration options, but most people will just want to select them all.



the recompiled kernel is actually running before going any further and wondering why nothing is working.

Getting started with iptables

As well as the kernel capabilities, we also need to obtain the user space program used to setup our firewall, *iptables*. This can be downloaded from www.netfilter.org or from our distribution's FTP site. Any distribution which comes with a 2.4 kernel should also include the *iptables* tools. A quick check to ensure that everything is happy is to run:

```
iptables -nL
Chain INPUT (policy ACCEPT)
target prot opt source destination
Chain FORWARD (policy ACCEPT)
target prot opt source destination
Chain OUTPUT (policy ACCEPT)
target prot opt source destination
```

If something different appears, ensure that the current release of *iptables* is the one being used, as older distributions will have older releases of the *iptables* tools, which may cause problems.

iptables splits the packet handling into three different tables, each of which contain a number of chains. The firewalling rules, which we create, are included within a particular chain. The three tables are filter, nat and mangle, all of which we will be looking at in this two-part *Linux Pro* series. The first is quite obvious, and is used for packet filtering. nat is used to provide packet modification capabilities, such as NAT/PAT and, of course, IP masquerading. The final table is the most obscure and is used for setting packet options, such as the Type Of Service (TOS) field and marking packets for further filtering or routing.

We're starting off by looking at the filter table, which contains three chains. These three chains allow us to filter specific traffic on our system. The INPUT chain is used for traffic which is entering our system and belongs to an IP address which is on our local machine. OUTPUT is used for traffic which originated on the local system, otherwise known as the firewall. Both of these are not used for traffic which is being routed between two network



Each chain in the filter table checks packets depending upon how they are routed through, to or from our firewall.

NETWORK SECURITY

interfaces on our firewall, which is left for the FORWARD chain. Using these three chains, we can easily block traffic we don't want, plus we can clearly define what type of traffic to allow onto our internal LAN.

The basic syntax of an *iptables* command is:

```
iptables -A INPUT -s 10.0.0.0/8 -j ACCEPT
```

This would add a rule into the INPUT chain, which matches any packet with a source address in the 10.0.0.0/8 netblock. If a packet matches this criteria, then it would use the ACCEPT target, which simply allows the packet on through. *iptables* supports CIDR and netmask syntax for defining networks within rules, so using 10.0.0.0/8 and 10.0.0.0/255.0.0.0 are equally acceptable. When adding a rule, it is added onto the end of the chain of rules, so will be the last one checked. As within the filter rule, whichever rule matches first will be the target for the packet and no other rules will be checked. We can also insert a rule into a specific location of the chain, ensuring that it is checked before other rules. Instead of using **-A** we use **-I** instead:

```
iptables -I INPUT 3 -s 10.0.0.0/8 -j ACCEPT
```

This would insert the rule into the third 'slot' of the chain. Of course, if there is currently only one rule in the chain, it would just be tacked on the end. For speed, if we neglect any slot value with **-I**, the rule will be inserted into the top of the chain, making it the first rule which is checked. Removing rules from a chain is equally as easy. We can either remove a rule in a specific slot location, or remove it based upon the options used when inserting or adding the chain.

To delete the first rule in the chain, we would do:

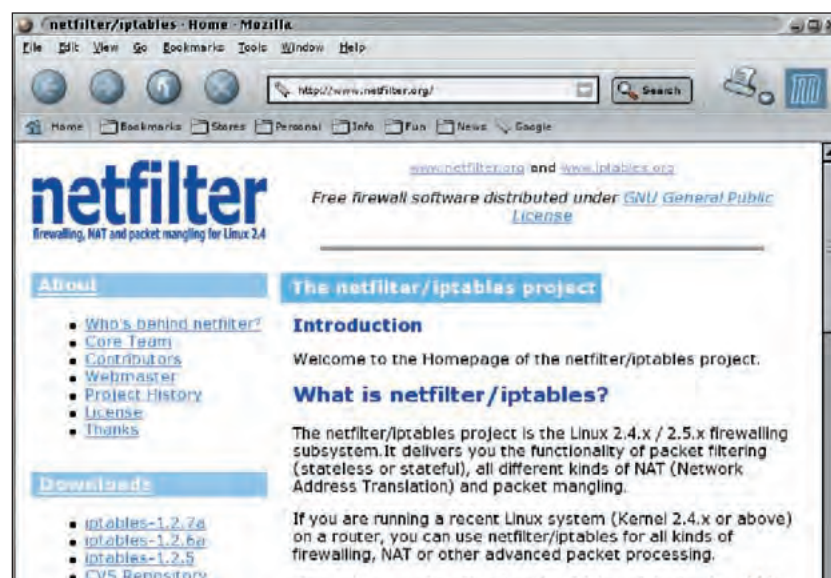
```
iptables -D INPUT 1
```

Or to delete the above rule which we inserted:

```
iptables -D INPUT -s 10.0.0.0/8 -j ACCEPT
```

If for some reason there are two rules within a chain with exactly the same options, then **-D** will delete the first one it finds in the chain.

There are three main targets for a rule within the filter table. The first we've already looked at, which is the ACCEPT rule. This allows the packet to be passed through the firewall without any noticeable interaction with the firewall. Of course, we want to block packets with our firewall, and *iptables* supports two different ways to do this. The first is the DROP target which simply drops the packet



www.iptables.org/ contains all the information on installing the user space utilities required to properly configure your firewall.

to the floor, as if our system was not even there. This type of packet blocking is known as 'stealth' firewalling, as it the firewall does not respond to the client when it blocks a packet. The other way we can block a packet is with the REJECT target, which drops the packet then sends a ICMP reply back to the client telling it why the connection failed. The default response is icmp-port-unreachable, although by using the **--reject-with** option we can send an alternative ICMP packet back, such as icmp-net-unreachable, icmp-host-prohibited, and in the case of a TCP connection, we can send a tcp-reset back which causes the client to reset the TCP connection to our system. Which ICMP packet we choose to respond with depends upon the impression of our network we want to supply to the remote system..

Matching packets

Matching packets is, of course, the most important aspect of our packet filtering setup, so we need to be able to clearly define which packets we want to block and which we want to allow through. The two most basic match conditions are source and destination address of the packet, the first of which we looked at earlier. These can either be individual IP addresses or a whole netblock, depending upon what we're trying to achieve. If we wanted to block packets heading to 192.168.1.2 from anything on the 10.0.0.0/8 network, we would do:

```
iptables -A INPUT -s 10.0.0.0/8 -d 192.168.1.2 -j DROP
```

We can also match based on protocol used, such as TCP, UDP, ICMP and so forth, as well as the specific port or service type used by that protocol. As an example, a common usage is to block connections to port 113 via TCP, which is used by identd:

```
iptables -A INPUT -p tcp --dport 113 -j REJECT --reject-with tcp-reset
```

Of course, we can mix the protocol and source or destination address into one whole rule, as appropriate.

```
iptables -I INPUT -p tcp --dport 113 -s 10.0.0.0/8 -j ACCEPT
```

When we specify a protocol we can either use the abbreviated name, such as **tcp**, **udp** or **icmp**, or we can use it's numeric reference, 6, 17 and 1 respectively. If for some reason *iptables* complains about a protocol name,

Creating a bash script containing firewall rules is the easiest way to maintain a system with huge rule sets.

```
david@macha:~ (pts/13)
iptables -A INPUT -s 127.0.0.1 -j ACCEPT
iptables -A INPUT -s 10.0.0.0/8 -j ACCEPT
#iptables -A INPUT -d 194.159.156.0/26 -j DROP
iptables -t nat -A PREROUTING -m state --state INVALID -j DROP

####
# Allow internal machines to access external gw IPs
####

iptables -t mangle -A PREROUTING -d 194.159.156.0/26 -s 10.0.0.0/8 -j MARK --set-mark 7

iptables -t mangle -A PREROUTING -s 127.0.0.1 -j ACCEPT
iptables -t mangle -A PREROUTING -d 10.0.0.0/8 -s 10.0.0.0/8 -j ACCEPT
iptables -t mangle -A OUTPUT -d 10.0.0.0/8 -j ACCEPT
iptables -t mangle -A OUTPUT -j MARK --set-mark 100

iptables -t nat -A POSTROUTING -p tcp --dport 3128 -d 10.1.1.4 -j SNAT --to 10.1.1.1
iptables -t nat -A POSTROUTING -m mark --mark 7 -j SNAT --to 10.1.1.1
iptables -t nat -A POSTROUTING -o lo -j ACCEPT
iptables -t nat -A POSTROUTING -m mark --mark 100 -j SNAT --to 194.159.156.145,118%
```


ensure that it is defined in `/etc/protocols`, as the system uses that file to associate protocol names with the numeric assignments. Ports are setup the same way, in `/etc/services`, although it's worth checking the port number against the service name to ensure that nothing has been changed. We don't want to block 'auth' expecting that to refer to '113/tcp,' when in fact someone has changed it to refer to '1113/tcp.'

We can also specify a 'match' option, using the `-m` flag. This allows us to use a kernel module to provide extra packet matching capabilities, the most popular usage of which is for connection tracking matching. By matching based upon the state of the actual connection, rather than simply the packet, we can ensure that connections which originated within our network are permitted, even if everything else is dropped. It can also prove useful for UDP packets, as UDP is a connectionless protocol, as it will track what we sent out and allow that system to send a reply to us through the firewall without opening it up.

The 'state' match has four different types of connection which we can match against. The most useful is `ESTABLISHED`, which corresponds to a connection which is already up and running. If the connection originated within our network, as soon as the packet passes through our firewall on its way to the Internet, it is tracked as `ESTABLISHED`. We also have a `RELATED` connection state, which is provided by a protocol helper module. The most common use for this is with FTP by using the `ip_conntrack_ftp.o` module, which allows us to track FTP connections back into our network properly, as when we download from a FTP server, it will try to make a TCP connection back to our system. There is also a 'NEW' state, which means that the packet is part of a new connection, meaning that it has not yet been tracked by the connection tracking system. Finally, we have the `INVALID` state which means that the connection is in an invalid state, so generally these should be dropped.

As a basic rule, we want to allow all `ESTABLISHED` and `RELATED` packets into our network, and selectively allow `NEW` packets through depending on the destination port.

```
iptables -A INPUT -m state --state INVALID -j DROP
```

```
iptables -A INPUT -m state --state NEW -j DROP
```

```
iptables -A INPUT -p tcp --dport 22 -m state --state NEW -j DROP
```

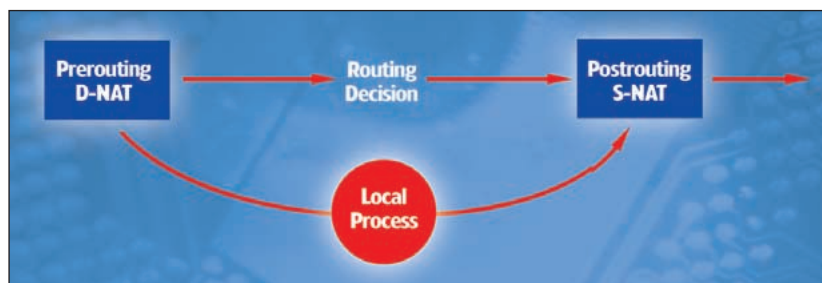
```
iptables -A INPUT -m state --state RELATED, ESTABLISHED -j ACCEPT
```

The above is a simple firewall configuration, which allows everything which originated within our network to be handled properly, plus it allows ssh in. However, we only want to drop `NEW` connections if they originated from the Internet, as we want to use other services internally, so we need a way to match based on the network connection the packet originated on.

This is done using the `-i` and `-o` flags, which refer to the 'in' and 'out' interfaces. The `INPUT` and `OUTPUT` chains can only use `-i` and `-o` respectively, but 'FORWARD' can use both `-i` and `-o`, as the packet is being passed between two network interfaces.

If we just want to block `NEW` packets which came from our PPP connection to the Internet, we could do:

```
iptables -A INPUT -i ppp0 -m state --state NEW -j DROP
```



The three NAT chains try to match packets at various stages of the kernel's routing decision for the packet.

or, if the PPP connection may have a name other than `ppp0`, we can drop `NEW` connections on all PPP interfaces:

```
iptables -A INPUT -i ppp+ -m state --state NEW -j DROP
```

When we're dropping packets, it's often important to make a log of these so they can be inspected for any attack against the network. To log packets which are matched by a particular rule, we need to set the target to `LOG`. If the rule already exists, we need to duplicate it and set the target of the first rule to `LOG`, as the `LOG` target will log the rule then allow the system to continue to search within the chain for another rule which matches.

```
iptables -A INPUT -i ppp+ -m state --state NEW -j LOG
```

```
iptables -A INPUT -i ppp+ -m state --state NEW -j DROP
```

Network Address Translation

NAT allows us to modify the source or destination address or port of a packet, allow it to be redirected, or so that it appears to come from another system. The most popular use of NAT is for IP Masquerading, where we have a LAN of systems using private IP addresses, which have network connectivity through a gateway which performs NAT on their packets so that all connections appear, at least to the outside world, to have come from the gateway system.

The `nat` table has three different chains, each of which is checked at a particular position in the routing of a packet. The first is `PREROUTING`, whose rules are checked for a match before the system attempts to route the packet anywhere. `PREROUTING` is where we perform destination NAT, or `DNAT`. We can also perform `DNAT` within the `OUTPUT` chain, which is for packets which originated locally, as they are never checked by `PREROUTING`. Finally we have `POSTROUTING`, which is the last chain checked for a rule match, which is where source NAT, or `SNAT` is performed.

With a static IP address, we can easily perform IP Masquerading using the `SNAT` target, by modifying all outgoing packets with a source address matching the external IP of our gateway.

```
# iptables -t nat -A POSTROUTING -o eth0 -j SNAT --to 194.159.156.1
```

For dynamic IPs, there is the 'MASQUERADE' target, which is a special type of `SNAT`. Not only does it mean we don't have to provide an IP address, but whenever the interface goes down, tracked masqueraded connections across that interface are dropped. For static IP addresses, if we use `masquerade`, if the connection goes down, when we come to reconnect our existing TCP connections will be dead, making the use of a static IP somewhat limited. We setup masquerading with:

```
iptables -t nat -A POSTROUTING -o eth0 -j MASQUERADE
```

Comments on this article to the usual address please. ■

NEXT MONTH

We're going to take a look at some more complex uses of NAT, the mangle table as well as GUIs aiding in the configuration of *iptables*.

More information, including HOWTOs and other documentation, on *iptables* and *netfilter* can be found at <http://iptables.org/>