

TURN OVER FOR YOUR FREE LINUX PRO MAGAZINE

STOP PRESS!
NOVELL BUYS SUSE
FOR \$210M

LINUX

THE UK'S BEST-SELLING LINUX MAGAZINE!

FORMAT

FREE YOUR OFFICE

Linux can deliver on the office desktop – find out how! **p44**



GETTING THE MOST OUT OF LINUX...

Every issue of Linux Format is filled with the best practical tutorials. This issue...

- MIGRATE FROM WINDOWS** p62
- HTML AUTHORIZING AND CSS** p66
- DTP WITH SCRIBUS** p70
- USING BLENDER SHADERS** p74
- ADDING DEPTH WITH GIMP** p78
- MASTERING PEAR AND PHP** p82
- REMOTE ADMINISTRATION** p86

KING OF GAMES!



First new release from Linux Games Publishing – run your own kingdom with a port of the PC classic Majesty!



SECURE THAT!

Novell's Steve Gaines rebuffs Microsoft's security claims **p12**

NEW ITANIUM ON TEST

Intel's 6MB monster powers HP's rx2600, reviewed inside **p22**

CD issue also available Printed in the UK
LXF48 CHRISTMAS 2003 £6.49
  1 3
 Media with passion 9 771473 968012
 (£6.49 outside UK and Republic of Ireland)

www.linuxformat.co.uk

7 pages of your problems solved • BitTorrent explained • KDE 3.2 revealed & Xmas fun inside!

Buying time

By the time you read this, I'm sure there will have been plenty more statements about Novell's purchase of SUSE. For Novell, the deal certainly seems to make sense – it has been trying to re-invent itself as a Linux organisation for some time now, and the purchase of a company with the capability of deploying a standard Linux distribution on a huge range of hardware is obviously a boost. For SUSE too, this will surely mean more investment and a chance to spread further than just Europe. But what is the bigger picture of the Linux market?

Red Hat and SUSE have been more-or-less concentrating on the enterprise market for some time. Mandrake are still committed to the ideals of Linux everywhere (especially the desktop), but lack the resources to push it further. Meanwhile on the sideline sit companies like Sun, reluctant to get into the business of actually developing a Linux distribution themselves, but keen to use them nevertheless. The Sun Java Desktop, which we'll be covering in more detail next issue, is

essentially Sun's own tools layered on top of SUSE. But, Sun points out, they could just as easily sit on top of Red Hat. Or maybe Mandrake?

It's quite likely that the spate of buyouts isn't over. There are still more technology companies with large amounts of cash, and plenty of smaller Linux outfits to choose from. The trick is choosing the right ones...

Meanwhile, our lead feature this month takes another look at Office software, particularly *Star Office* and *OpenOffice.org*, thanks to their new releases. Is there any progress being made towards apps that are credible replacements for the ubiquitous market-leader? Find out on page 44.

Microsoft's Steve Ballmer recently made further efforts to either a) undermine the credibility of Linux security; or b) reinforce the notion that he doesn't understand Open Source development. Take your pick, but we have comment from Novell on security, p12.

And there's KDE. And BitTorrent. And more!
A very Happy Christmas to all from us here at LXF!



Nick Veitch EDITOR

An office free from MS Office? Wake up! It isn't just a dream after all! **p44**

Apparently Windows is more secure than Linux because MS don't issue as many patches. Hmm... **p12**

A sneak preview of the new features of KDE 3.2 **p50**



12



50



44

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...



Andrew Channelle
The Linux beginners' best friend, Andy also dashed off the feature about Novell and SUSE as our presses rolled...



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



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



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



Richard Cobbett
Part-time penguin Cobbo is Features Ed on LXF's sister-mag *PC Plus*, and author of this month's Office cover feature.



Richard Drummond
Rich lives in a different time zone these days. GMT-6. Weeks, that is.



Hoyt Duff
Fishing pier proprietor Hoyt is also co-author of the *Red Hat Linux 9 Unleashed* book and a Mandrake contributor.



Neil Bothwick
Sources the very best software for our discs. Just Googling his name turns up more than a thousand references...



Michael J Hammel
Professional GIMP artist who pens (or pencils) our current Open Source graphics tour-de-force.



Paul Hudson
Our Reviews Ed has been to Hungary and China in the same month – maybe he's moonlighting as an air-stewardess?

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

LXF48 Christmas 2003

Contents

LINUX
 FORMAT

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

COVER HIGHLIGHTS

FREE YOUR OFFICE!

44

 With *OpenOffice.org* and *StarOffice*, Linux can deliver on the office desktop!

12 Secure THAT!

Novell's Steve Gaines rebuffs Microsoft's Linux FUD

22 HP Itanium server on test

Integrity rx2600 powered by the Intel 6MB monster

26 LINUX GAME: Majesty

Usurp the throne and start your own reign of terror!

70 DTP with Scribus tutorial

 Who needs costly *InDesign* or *QuarkXPress* licenses?


'Whilst **close integration** means **Microsoft** can provide rich functionality, at the same time this **opens up security flaws.**'

12 Steve Gaines of Novell testifies for Linux

'In a world that seems to be at one another's throats, it's nice to be a part of something fundamentally good...'

07 Jono Bacon climbs the ethical tree

LINUXPRO

Security The risks for sysadmins, as seen by ISS's CTO **SCO vs Linux** Implications for UK **Java Development** JBoss marches on **Enterprise** Novell on Open Source **Hosting** Host Europe targets SMEs and more



22 This or a BMW? Our very first high-end (and high priced) Itanium server.



6 News

STOP PRESS: Novell buys SUSE!

14 Mailserver and Reader Tips

Got an opinion? Have your say!

20 Slackware 9.1 review

Latest version of Linux's oldest distro

24 Gigabyte SR125ED server review

Does it compare to its predecessors?

28 Book reviews

Two Red Hat giants head-to-head

32 Roundup: MIDI sequencers

Sound applications rested & rated

38 HotPicks

All the latest Open Source software

50 KDE 3.2 — what the future holds

Our predictions and comments

56 What On Earth... is BitTorrent?

Lightning-fast downloads ahoy!

62 Beginners' tutorial

Migrating to Linux

66 HTML/CSS tutorial

NEW SERIES! Linux and web code

70 Scribus tutorial

Free desktop publishing is here...

74 Blender tutorial

Creating an atmosphere

78 The GIMP tutorial

Adding depth with shadows

82 Practical PHP tutorial

What a lovely PEAR!

86 Server School tutorial

Remote access: how to do it

92 Answers

Your queries answered

97 Backissues and subscriptions

Great new money-saving offers!

102 User Groups

Local support and fun

104 Happy LXF-mas!

Seasonal puzzles and silliness!



COVERDISCS

A DVD or 3 CDs packed full of the latest Linux goodies **106**



CDS A, B AND C

SUSE Live Eval 9.0, Slackware 9.1 BOTH boot directly from CD; **gentoo** Flexible and full-featured file manager; **Karchiver** one-click tarball and zip archive handling; **BitTorrent** Download and share large files over any connection; **Mozilla, Opera** and **Firebird** up to date, standards-compliant browsing; **Mambo Open Source Site Server** Create and manage any size of web site with ease



DVD

KDE 3.2 alpha 2 see what people are raving about with this preview; **DamnSmallLinux** distro; **Audacity** professional-quality sound editing; **Samba 3.0** the latest Windows sharing app

Please read the coverdisc instructions on page 109 before installing from coverdiscs!



12

50

26

56

SAVE MONEY!

SUBSCRIBE TO LINUX FORMAT

See page 98 or phone 0870 4448645



Newsdesk

- Mandrake 9.2 eats LG CD drives ● Lightning LINUX ● SCO News ● W3C's Tim Berners-Lee appeals
- Open Government UK? ● CrossOver Office gets Flash ● BarbieOS ● Red Hat migration to Fedora

\$210MILLION DEAL CREATES NEW FORCE FOR ENTERPRISE LINUX

Novell – desperately seeking SUSE

SUSE, Europe's biggest Linux vendor, has been acquired by Novell Inc. In a long rumoured move, Novell has signalled its intention to snap up the German company for US\$210 million in cash. Novell's share price has soared by over 20 per cent after the announcement was made.

Jack Messman, Chief Executive, was in SUSE's German headquarters to announce the move. He said it was "a very exciting day for Novell, SUSE and the Open Source community." The acquisition, he maintained, was good for all parties and would create the first real opportunity for lowering the barriers to the adoption of Linux across the enterprise.

"We are building a trusted partnership with a unique expertise," he said. "Together we will work to bring Linux from the periphery of the network"

For Novell, the benefits are obvious. After the announcement earlier this year that Novell services would run on either the NetWare or Linux kernel, and the acquisition of Ximian – which provided the very useful *Evolution* and *Connector for Exchange* combination – Novell has added the "final piece of the jigsaw" and now has the "entire enterprise covered from the server to the desktop".

"This acquisition will create the first – only – \$1 billion Linux company able to provide a full end-to-end stack and,

more importantly, the global infrastructure capable of supporting it." But it would also benefit SUSE, providing a foothold in the USA where traditionally Red Hat dominates; and new and existing SUSE customers, who would gain access to the "Global enterprise level support that Novell is known for".

Messman added that the SUSE division of Novell would also keep its presence in Nuremberg and would not scale back in its support of ground-

"We are building a trusted partnership with a unique expertise, together we will work to bring Linux from the periphery."

breaking Open Source applications. "We expect to learn a lot from them," he said.

Smooth transition

A survey of staffing levels didn't reveal a significant overlap of skills, and most of the SUSE development will be staying with the organisation to ensure a smooth transition from an independent operator to a division of a larger corporation.

Richard Seibt, SUSE CEO, said that the vendor's absorption into Novell would help achieve the objectives SUSE had always set itself. "Our goal has always been to make our

technology available across the world, to become the *de facto* standard for Linux." He said the addition of established support channels would also improve the standing of Linux among technology buyers. "We can drive the adoption of Linux much faster together than we ever could apart"

Some analysts have pointed out Novell's record of 'jumping on the bandwagon' before leaping off just as quickly when something new comes



Jack Messman, CEO Novell

the company by IBM, already a big backer of SUSE in Europe. A move likely to add fuel to the fire in the SCO case.

Jack Messman said there was often a "lot of fluff" in these announcements, "but when people are putting in real money, you know they mean business."

"This is an important moment in the history of the Linux marketplace."

With regards to SCO, Messman said the company's position had not really changed and the case would not affect the foreseeable future for Linux in general or Novell/SUSE in particular. "We're still calling for full disclosure, but we are moving forward aggressively to support Linux despite these unsubstantiated claims."

There were many reasons, Messman said, why Novell was committing itself so wholeheartedly to this radical move, but the simplest one was that "Linux is the future of computing."

along, often leaving a dead product and/or a lot of work for lawyers. The UNIX source code, Word Perfect, Java for NetWare...

Chris Stone said the strategy Novell has embarked on, and put so much effort into, was not without its risks, but that it was vital to do it right. "We are transitioning a worldwide corporation from a proprietary software house to an open Linux operation. We did not want to come in as a 'big bad' enterprise company. We wanted to prove we were willing to participate."

Novell's senior management team also used the occasion to announce an investment of some US\$50 million in

LINUX HOSTING

Server host seeks new partners

Rackspace, which offers managed server hosting on both Red Hat Linux and MS Windows Server, has expanded its successful customer referral program and launched a Worldwide Business Partner Program. The new initiative aims to build on the referral program which Glenn Reinus, Vice President of Sales, says brings in more than 50 per cent of new customers.

"To show our commitment to those partners, we're launching a worldwide program that incentivises partners to

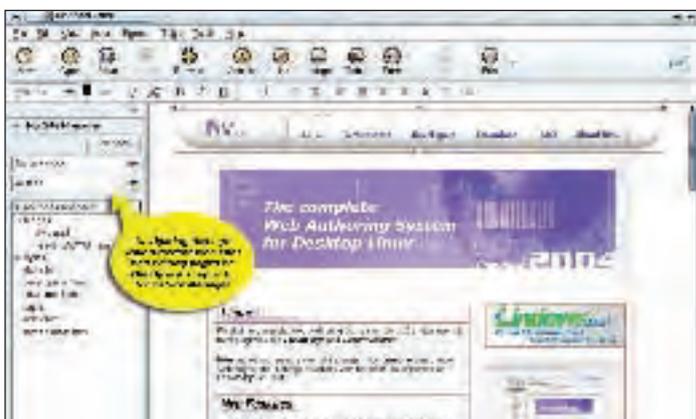
extend Rackspace's service offering to their customers. We've increased the level of sales and marketing tools we provide our partners, better supporting their efforts to grow revenues and build their businesses."

Enterprises that join the Business Partner Program gain access to Rackspace's exclusive partner portal, which contains a range of sales and marketing tools, and will be included in the company's Online Directory.

www.rackspace.co.uk

WEBSITE EDITOR

Plugging the WYSIWYG gap



Nvu expands the capabilities of the Mozilla Composer.

Michael Roberts, CEO of LindowsOS and Linux publicist extraordinaire, has noticed a gap in the application market and has set about filling it with *Nvu*, a *Mozilla Composer*-based website editor. Moreover, in a very Open Source-friendly move, Lindows has hired the current *Composer* project leader Daniel Glazman to expand the application, often the neglected triplet in the *Mozilla* family, into a user-friendly editor aimed at *Frontpage* and *Dreamweaver* designers. The result will not just be limited to *Click 'N' Run* Lindows users, as it will be covered by *Mozilla.org*'s MPL license.

Roberts says the initial 1.0 release, which adds file management facilities and improved table and form tools, will be available in early 2004. Funds have

been set aside, he said, not just to get the application to release standard, but to continue development to create what he described as the 'premier Web Authoring System for Linux.'

In his monthly address, Roberts said the company would be committing money, engineers and ideas, but that Open Source developers would still be able to make major contributions to the project.

www.nvu.com/



The ability to manage sites from within the application is regarded as a vital feature.

NEWSBYTES

■ The first development release of **GTK+2.4**, featuring GTK+2.3.0, *Pango 1.3* and *Glib 2.3*, has been released. The new package includes some significant new widgets such as a unified *GtkComboBox* (to replace both *GtkOptionMenu* and *GtkCombo*), and a replacement for *GtkFileSelector*.

■ **Andrew Tridgell**, a leading light in Samba development, has been crowned Australia's smartest man (in the IT sector) by *Bulletin* magazine. "It's a great honour and it's really an honour for all the people who work in free software and Open Source in Australia," he said.

■ **Sun** has signed a four-year Whole of Government contract with the authorities in New South Wales, Australia, giving computer users access to a range of Sun applications including *StarOffice* and the imminent *Java Desktop*.

■ Desktop Linux vendor **Xandros** has appointed a new CEO. Andreas Tupaldos has worked in the computer industry for 30 years and was founder of *AxsOne*.

■ Low-end box shifter **CPU Builders** has added an easy to use hardware and software configuration tool to its bespoke Linux distribution. Aimed at Windows users, the new control panel handles system set up and online software installation and updates. It will be included in all their pre-built systems.



■ **Oracle** has had its database product certified for use on Red Hat Enterprise Linux 3. Certification, which involves over 1,000 rigorous tests is thought to be vital if Red Hat is to make headway in the Enterprise sector. This is the latest example of Oracle's Linux strategy which, earlier in the year saw the company launch a \$100 million campaign to encourage the adoption of the OS among Oracle application developers.

■ Douglas Steele has written a paper on the legal aspects of projects such as **XPde**, which is attempting to build a Windows-like desktop for Linux, when it comes to issues of intellectual property and copyright. The document would make interesting reading for any UI developer, but it also works as a readable history of the development of graphical user interfaces. Download it from www.xpde.com/

Jono Bacon

The founder of UK Linux, *KDE* developer and all-round nice guy, Jono is also a musician who's tunes have been featured on *Slashdot*.



COMMENT

Important ethics?

“ At the moment, we seem to living in rather troubled times – not only the political issues going on in the world, but also the issues with the SCO vs everyone + dog case rumbling in the tech headlines. In times of need, the typical human response is to club together and fight the threat as one. This has been seen in times of war, disaster and other negative situations. In our current times, the SCO case not only presents a challenge the very fabric of our community, but also highlights some of the issues that we have perhaps glossed over in the past. Our community is facing a challenge that involves each aspect of it being examined and tested.

One aspect of our community that I am proud of is the way we have worked together to fill the gaps in our OS. Each area has particular developers working on particular issues, and although these scattered contributors are facing their own issues, there is one aspect that binds them together – the vision of a free OS and tools for everyone.

The thing I find compelling about this vision is that ethics is firmly seated at the forefront of the effort. Equality is the key here, and in a world where everyone seems to be at everyone's throats, it is nice to be part of a community that is doing something fundamentally good. I have nothing against companies making money with our OS, but the basic ethical fabric of the community needs to remain. Hopefully the SCO case will work out for the better and allow free software to thrive alongside businesses. I firmly believe these two disparate cultures can work together. ”

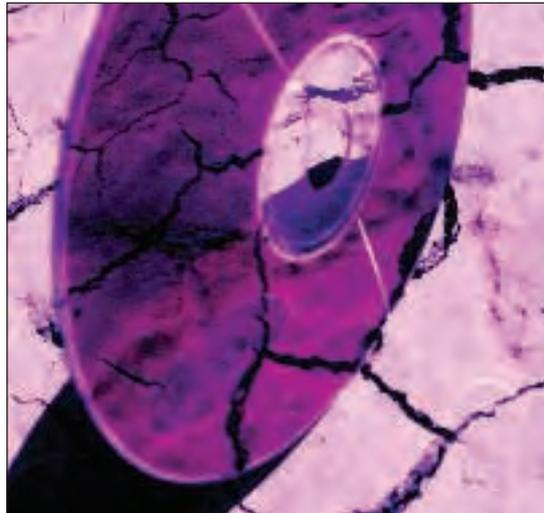
BLAMESTORMING

Mandrake 9.2 CD problems

A major hardware problem has been discovered in the latest distribution release from Mandrakesoft, which was also featured on last month's *Linux Format* cover disc. The kernel included with Mandrake 9.2 has been shown to cripple some CD-ROM drives made by LG Electronics –

included in machines from Dell, Compaq, Hewlett Packard and others. Users with an affected system who try to install Mandrake 9.2 will first be told that the base system could not be installed, and then rebooting would reveal the CD drive to be dead. The difficulty has arisen due – depending on your perspective – either to LG's

LG and Mandrake – living proof of why standards are so important!



SOME AFFECTED DEVICES

CD-ROM drives with the following device numbers have been reported not to work:

GRC-8521B	GCD-R560B
CRD-8522B	GCD-R542B
CRD-8521B	GCD-R540C
CRD-8520B	GCD-R540B
CRD-8483B	GCD-R520B
CRD-8480C	GCD-R420B
CRD-8480B	GCD-R400B
CRD-8400C	GCD-R320B
CRD-8400B	GCD-R300B 4
CRD-8322B	GCD-R200B
CRD-8320B	
CRD-8240B	
CRD-8241B	
CRD-8160B	
CRD-8161B	
CRN-8240E	
GCD-R580B	

The following have been tested by various users and should pose no problem:

LG GMA-4020B DVD-RW
 LG GCC-4120B CDRW/DVD
 LG CD-RW CED-8080B DVD/CD-R/RW/CDROM
 LG CD-RW CED-8120B
 HL-DT-ST GCC-4480B (firmware 1.01. Firmware 1.00 is reported to have the problem)
 HL-DT-ST RW/DVD GCC-4480B
 HL-DT-STDVD-ROM GDR8161B
 HL-DT-ST CD-RW GCE-8240B
 HL-DT-ST GCE-8481B CD-CDRW
 HL-DT-ST CD-ROM GCR-8520B
 HL-DT-ST GCE-8520B

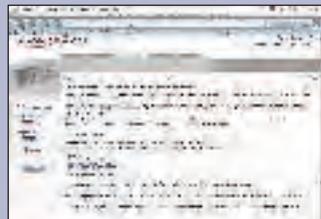
To discover your device number, open a console and type `dmesg|grep hdc`, where `hdc` is the name of your CD drive.

WINDOWS DRIVERS FOR LINUX

Canadian support consultant Linuxant has developed *DriverLoader*, an application which, the company claims, works as a 'revolutionary compatibility wrapper' that enables Linux users to use standard Network Driver Interface Specification (NDIS) drivers that commonly ship with Windows hardware. The big feature of the latest release is multiple driver support and compatibility with IntelPRO (Centrino), Intersil Prism GT, Broadcom and Atheros-based wireless systems.

Linuxant says it intends to widen support for 'the majority of Windows NDIS drivers' in the future and hopes, with support of hardware vendors, to keep end-users licenses available on a cost-free basis.

All Linux-specific code in the product is Open Source and can be used under any supported version of the kernel. Binary packages are available for Red Hat, SUSE, Mandrake and Debian based on 2.4 or 2.6 kernel. www.linuxant.com/driverloader/



Driverloader now has support for a wide range of Windows only WLAN cards.

unorthodox use of the FLUSH_CACHE command, which causes these drives to erase their firmware, or the 'packet writing' patch which was introduced into to Mandrake's 2.4.22-rc2q5 kernel on August 15th.

Mandrake developers have stressed that it is mainly basic CD-ROM drives which are affected, not DVD-ROM or the majority of CDRW drives, though this depends on their firmware version.

In an alert, Mandrakesoft put the blame firmly on the shoulders of LG Electronics, claiming that LG's "drives

are not compliant with the ATAPI specification. The specification does not require an implementation of the FLUSH_CACHE command in the driver, and returning an error (or doing nothing) would have been the correct behaviour for the drive."

To remedy the situation, a new kernel has been released (2.4.22-21mdk) and CDs and ISOs should be available as you read this. The most up-to-date list of devices is available at <http://mandrake.vmlinux.ca/bin/view/Main/LgCdrrom>

Embedded Linux News

● **Aplware** of Switzerland has announced the latest version of its Lightning LINUX distribution which offers a useful package of communications and networking tools, high-availability clustering tools and a complete remote management system. Lightning LINUX 3.6 is designed for integration into networking gear such as routers, firewalls and VPN gateways, but is also suitable for implementation in IP PABX products. www.aplware.ch/

● **Cadence Design Systems** has announced a Linux port of its IC packaging and PCB design solutions to Linux, a first for the platform. Version 15.1 of the suite of products includes PCB Design Studio, which was formerly a Windows exclusive application, and PBC Design Expert. www.cadence.com

● Processor builder **ARM** has become the latest big name to join the Consumer Entertainment Linux Forum (CELF). ARM spokesperson Mary Inglis said the company could play a vital role in making the platform viable for a whole range of consumer devices. 'Changes from both a technical and architectural standpoint are needed to help the platform reach its full potential in a consumer electronics environment', she said. www.celinuxforum.org/



● **Trintech** is riding the wave of changes taking place in the credit card industry with the launch of the first Linux-based 'Chip and Pin' device for secure card transactions. Over the next few years, devices such as the Payware Smart5000 will appear in place of the usual swipe and sign systems which credit card companies claim cost the industry billions of pounds through fraud. The device is built on Trintech's OpenPay system which marries a small, fast ARM processor with a highly tailored version of Linux. It recently won a transaction speed competition organised by the German banking industry. www.trintech.com

SCO News

■ SCO's Invoice plan – by which the company sends out a selection of invoices to prominent Linux users demanding payment – was launched once again, before being shelved and then resurrected with the caveat that they would only be invoicing Fortune 1000 corporations. They also extended the deadline for license buyers to qualify for a reduced, bargain basement price of US\$699 per CPU. Fears that sending out invoices may (allegedly) embroil SCO executives in mail fraud charges if the impending lawsuit fails are said not to be behind the move.

■ SCO's legal strategy is becoming clearer. The intention appears to be to have the GPL declared void on the grounds that it is unconstitutional.

■ Meanwhile in India, Life Insurance Corporation of India (LIC), the largest insurer in the country, has begun the process of migrating 6,000 UNIXware servers across 2,048 branches to Red Hat. A spokesman told *The Financial Express* that they had looked at Windows but opted for UNIXware because they didn't want to be tied to a 'license-based OS'. LIC appears on SCO 'success stories' web page.

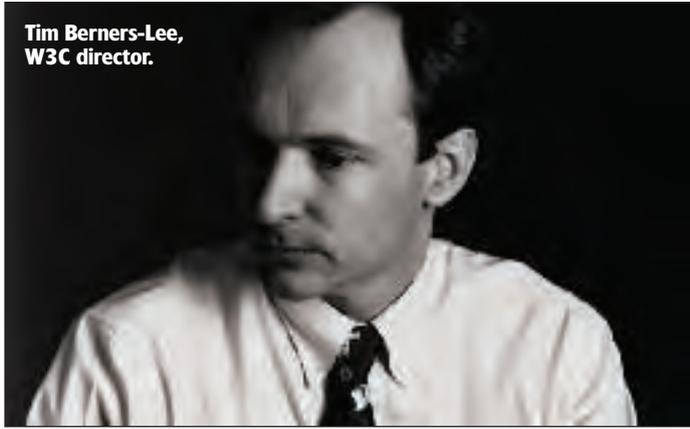
■ In an interview, Darl McBride has suggested that IBM would be in hock to SCO to the tune of \$50 billion per year. Revenue from Linux licenses would also, he said, bring in billions of dollars each year.

■ As ever, the dedicated people at Groklaw (www.groklaw.net) is documenting the case in the detail that it deserves. Also see this month's *Linux Pro* for IP barrister David Harris's take on the whole contra-temps.

WEB PATENTS

IE plugin battle escalates

Tim Berners-Lee, W3C director.



The director of the World Wide Web consortium, Tim Berners-Lee, has made an unprecedented appeal to the United States Patent and Trademark Office to re-examine patent number 5,838,906 in the light of a number of claims of 'prior art'. He also highlighted the impact that upholding the ruling would have on the future of the web and access to the vast body of work that already exists. The patent was at the centre of a recent lawsuit brought by Eolas, an intellectual property broker, against Microsoft's *Internet Explorer*, though the result could impact significantly on all current web browsers.

In his letter to James E Rogan, Berners-Lee said the patent, granted in 1998, could potentially "undo years of work that have gone into building the web".

Examples of prior art, some of which were not examined in the original trial, include the EMBED tag featured in two published papers by HP researcher Dave Raggett that works in an identical fashion to the EMBED tag in the patent.

In response to the letter, the US Patent Office may order an official re-evaluation of the patent, which could take a further 90 days. In the meantime, Eolas are attempting to have shipments of Internet Explorer halted while Microsoft's appeal works its way through the legal system.

Other examples of prior art include the *Write* word-processor that shipped in Windows 3.1 and could embed images from *MS Paint*, the NCSA *Mosaic* web browser – both cited in Berners-Lee's letter – and Ray Ozzie's demo of similar processes in *Lotus Notes 3*.

LINUX TCO ON TEST

Open Government in UK?

Recently, IBM has won a contract to begin trials of Open Source software across a number of UK Government departments. The nine trials have been initiated by the Office of Government Commerce (OGC), the body set up to ensure government got value for money in the public services, and the Office of the e-Envoy. www.e-envoy.gov.uk/Home/Homepage/fs/en

The OGC began to look seriously at Open Source alternatives to its existing proprietary systems in response to the

experiences of the German federal government, in particular, the high-profile adoption of Linux in the city of Munich.

Martin Day, a spokesman for the OGC told Silicon.com that the study would look at value for money and, crucially, interoperability with legacy systems.

In response, Microsoft offered Newham Council, which had already committed to investigating the advantages of Open Source software, a fully funded audit to demonstrate how sticking with Microsoft products would provide a lower total cost of ownership

(TCO). The audit, to be carried out by Cap Gemini Ernst and Young, will be funded by what UK tech news site The Register www.theregister.co.uk calls Microsoft's 'Stop Linux Slush Fund'.

The big danger for Microsoft is that, according to Ted Schadler, analyst at Forrester Research, business users see government adoption of Open Source technologies as a legitimising force.

Let's hope that contact with the Open Source ethos persuades UK government that browser-agnosticism for its sites is a priority issue too...

Hoyt Duff

The co-author of *Red Hat Linux 9 Unleashed* runs a fishing pier when he's not shouting lovingly at his computers.



COMMENT

HFG or LFG?

“ Do you only enjoy activities revolving around your computer? Prefer solitary activities? Preoccupied with fantasy or introspection? Have few close relationships? Flaunt convention in behaviour or costume? Flaunt yourself in a costume *at* a convention? Display little interest in sexual experiences with another person? Congratulations, mate – you're a Geek.

To be more precise, you may have schizoid personality disorder, a condition that seems to describe many of us. I don't believe it's much of a problem; many of us get along with life pretty well and take pride in an eccentric lifestyle.

Why bring this up? On the Linux Lunacy Cruise in September 2003, it was interesting to observe the leadership of the Linux community as well as the other hundred or so attendees. Many of the leadership (and a few of the attendees) fell into what I call the High Functioning Geek category. While they exhibited typical Geek behaviour (crouched over a laptop in the coffee bar every day), they were also 1) married, and 2) capable of carrying on an articulate conversation that did not involve Linux or computers. The Low Functioning Geeks (there were a few extreme examples of schizoid personality disorder) were in awe of the HFGs and it's easy to see why the HFGs had become elevated to positions of leadership in the community.

Being an HFG has nothing to do with intelligence and everything with social and communication skills. All of these skills can be learned. In fact, that's essentially what the therapy for schizoid personality disorder is – I know this because a man I met in a pub who claimed to be a doctor told me so...

CODEWEAVERS GROWTH

Bigger CrossOver Office

The latest update of CodeWeavers' CrossOver Office adds much requested support for *Dreamweaver MX* and *Flash MX*, Macromedia's flagship web design applications. CodeWeavers' Jeremy White said *CrossOver Office* had become an important part of the push to put Linux on the corporate desktop with support for applications such as *Office XP*, *Lotus Notes* and *Quicken*. 'Now, with additional support for *Macromedia Dreamweaver MX* and *Flash MX*, combined with our existing support for *Adobe Photoshop*, *CrossOver Office 2.1* gives web developers and design firms the first Linux solution for their most important design applications.'

Quite unusually for this situation, Macromedia has come out in support of the project. Susan Morrow, the company's vice president of product management, said they were committed to supporting customers regardless of

the platform they choose to use. "We are excited to see how Linux developers will use *CrossOver Office 2.1* with its support for *Dreamweaver and Flash*," she said.

Like Adobe before it, industry experts claim, Macromedia has been under some pressure to port their applications to Linux, but support from *CrossOver Office* is seen by many to be a double-edged sword, offering developers a way of exploiting a growing market without the cost of creating a native port. Jeremy White, CodeWeavers CEO, says it just means software vendors get an opportunity to test the Linux waters without committing to an enormous and costly development effort.

"Frankly, the Linux market, while full of many passionate and amazing people, has not really proven itself to be rich ground for software companies, particularly those focused on desktop applications." It is, he says, the standard chicken-and-egg argument. "Folks can't move to Linux without their applications, and the app makers won't move without

a clear and large market. We're trying to break the egg and enable at least some migration, which will then hopefully lead to a larger market"

Fortunately in many cases, CodeWeavers do receive assistance from the original developers. "The team at Macromedia has been fantastic; we have been impressed by their support for our efforts as well as their practical and positive view of the Linux community," White said.

Support for Dreamweaver and Flash is flagged at the Bronze level which – officially at least – suggests the application isn't ready for prime time. However, White says the Bronze designation is an acknowledgement of the problems some users can have with *Wine* and that, as such all applications start at Bronze level. In response to a comment on this issue at Slashdot, White wrote: "We try to help our customers be cautious in their



adoption. For example, *Photoshop*, which actually is one of the very best performing applications in *CrossOver*, started at Bronze, and is now only at Silver. With that said, we have found *Dreamweaver* to be very complete, with only a few remaining bugs. And we have yet to find a bug in *Flash*."

Linux Web Watch/



Sugar-coated Linux hoax.



Share your money with *Ardour*.



Share your time with *GnuCash*.



Share your experiences at LDP.

Christmas is coming...

And the geeks are getting fat, please put a dollar in the pizza delivery boy's hat!

What do you get the girl who has everything? Her own OS of course! News of a new Linux-based BarbieOS built to grace the next generation of Mattel Barbie laptops were sadly revealed as a hoax. The site at <http://qrxx.4t.com/barbieOS.htm>, is quite convincing and fooled a few posters on tech sites and forums. Now for something serious. While

many Open Source projects garner lots of publicity and support from enormously rich hardware companies, there are some that subsist on good will and the part-time efforts of dedicated developers. Here's a few that might benefit from some user generosity this Christmas...

Ardour (<http://ardour.sourceforge.net/>) is a fantastic

musicians' tool that has been working towards a stable release for three years. It's nearly there and the developers have a PayPal account which could do with some attention.

GnuCash, the GNU way to manage your money, is more interested in donations of time, especially from those able to help document the applications extensive APIs or get

involved in the interoperability effort. www.gnucash.org

If you've discovered how to do something and want to share your knowledge with the world, you could do worse than writing a guide to add to the Linux Documentation Project. There's a decent authors guide at www.tldp.org/LDP/LDP-Author-Guide/index.html.

NEWSBYTES

■ **IBM** is offering cost incentives for purchasers of its revamped pSeries 615 server to choose Linux as the OS instead of AIX. The new server ships with IBM's dual-core 1.45GHz Power 4+ processor and Linux buyers will save US\$1,000 on the AIX price for a basic system. A two-way system running AIX will retail at US\$12,500 while the same machine with Linux will come in at US\$10,000. Analysts suggest that the move demonstrates the faith IBM is putting in Linux and will also fuel competition for the lower end of the market currently dominated by Intel-based solutions.

■ **Hewlett-Packard** has begun a campaign to lure Sun customers away from Solaris. The two corporations have a long-running enmity, but this strategy is different as HP's gambit is to play on the strengths of Linux instead of Windows or HP-UX. A HP spokesman said they were targeting customers who had yet to experience the benefits of Open Source.

■ Meanwhile **Sun** is attempting to lure low-end users to Linux with its US\$100 per-user per-year offer on its Java Desktop product.

■ **Swinburne University of Technology's** Centre for Astrophysics and Supercomputing has cut its computing costs by adding six Apple Xserve RAID machines to its 130 node Xeon/Linux cluster. The 13.2TB storage system came in at 'around \$100,000'. Professor Matthew Bailes told LinuxWorld that after some initial scepticism from Apple, the University had the RAID up and running 15 minutes after a kernel recompile.

■ **New Zealand's most prominent Mandrake user**, Green MP **Nandor Tanzcos**, is pushing for the national government to adopt a more open IT strategy. Tanzcos says the NZ government has joined Microsoft's Government Security Program, which allows limited access to the Windows source code, but there's no clear goal of the project. "If its going to cost X million dollars to participate in this scheme, why don't we just spend that money moving to Linux?"

■ **Richard Stallman** fell and broke his arm on some ice in Finland where he was due to give a speech at a seminar on copyright and free software. RMS eventually gave the talk over the phone from his hospital bed, before offering a ride to Helsinki for anyone conversant in *Emacs*.

■ **MySQL AB**, the commercial entity behind the Open Source database, has acquired Alzato, makers of NDB Cluster, in a bid to improve the scalability of the database. NDB Cluster, a data management system aimed at the telecoms and IP industry will be integrated into the next major MySQL release.

Red Hat Linux Migration Resource Center will help network customers to move to Enterprise Linux v3, and products like RHAS will continue to be supported – for a price, of course.



DISTRO NEWS

Red Hat boldly goes... Enterprise only

Red Hat Network officially told its customers in an email on Monday 3rd November that it "will discontinue maintenance and errata support for Red Hat Linux 9 as of April 30, 2004"; furthermore "Red Hat will discontinue maintenance and errata support for Red Hat Linux 7.1, 7.2, 7.3 and 8.0 as of December 31, 2003. Red Hat does not plan to release another product in the Red Hat Linux line."

This development should come as no surprise to regular *LXF* readers, as this story was first mooted in our News section a couple of months ago, but less-prepared Red Hat users seem shocked by this development. Some feel betrayed, expressing the belief that 'dead rat Linux' might be cutting off its nose to spite its face – a mood particularly prevalent among the 'old school' RH contributor community and small operations involved with selling, administering and advocating the benefits of FOSS to SMEs. The announcement is tempered a little by some hefty price cuts – if you purchase Red Hat Enterprise Linux WS or ES Basic before February 28, 2004, you will receive two years' access to updates, errata, security patches and new releases via RHN for the price of one year, for up to 10 units.

Red Hat's second-quarter 2003 profits of US\$3.3million and close cooperation with IBM are seen by most pundits as the motivation for the decision to concentrate on the enterprise market. So, from the end of this year, users wanting enterprise-level Linux with the Red Hat name on it will have to pay for it. To keep using a Red Hat-style distro for free, you can use Fedora.

<http://fedora.redhat.com/> maintains that: "The Red Hat engineering team will continue to participate in the building of Fedora Core and will invite and encourage more outside participation than was possible in Red Hat Linux. By using this more open process, we hope to provide an operating system that uses free software development practices and is more appealing to the open source community."

Neither of these options appeal to you? There are plenty of other Linux distros out there to choose from. Contrary to what the nay-sayers opine, Red Hat's decision is actually going to help other Linux organisations that are committed to desktop distributions – get links to a Linux distribution that suits you from the comprehensive list at <http://old.lwn.net/Distributions/>

David Cartwright

An IT consultant since before Turing's time, David specialises in Linux systems and solutions provision.



COMMENT

In my CUPS

“ The more I play with Linux, the more I find it irresistible as a mid-level server for small and medium enterprises (SMEs). When I sysadmined Sun-3s (in the days before SunOS was called Solaris), we ran our machines as AppleShare file servers, NNTP-based news repositories, email servers and the like – all of which were a faff to set up – you'd spend hours trying to figure out how to patch the AppleTalk protocol into the kernel, or set up the NN newsreader for command-line users.

The one thing I'd not set up for an SME clients was Linux- or Unix-based print serving. Not because I can't do it, but because instead of taking an hour or two to get my head around Unix-style printing (hellish in the old days) it's been quicker (= cheaper for the client) to simply turn on Windows printer sharing on someone's PC and stuff the printer on that.

Until this weekend, that is, when I was asked by someone whether they can connect a printer to their office-based Linux server and print to it over a VPN from their WinXP laptop. Maybe I'm just weird, but I find anything to do with Windows file and printer sharing a pain to do over VPNs (NetBIOS isn't great over a WAN, even if your VPN software is pretending it's a LAN) so I figured it was time I tried this Internet Printing Protocol thingy.

30 mins later, after a process that was 90% GUI-based and only 10% editing cups.conf, he was printing over the VPN from his PC, and I was wondering whether I'd missed something, as it all seemed too easy. Then I realised that this is the way Unixes, and more so Linux, is going: though it's getting more powerful (and a bit sprawling in some cases?) it's also getting easier. This stuff really is available to the non-experts. ”

Mailserv

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

Linux racism?

I just bought a copy of *LXF*. Generally I liked the magazine – it has a fairly good mix of articles for programmers, developers, novices and creative applications users. It is quite well written, if a bit self-congratulatory (though this is expected).

However, I was disturbed to see that all the images of people were of white people, except for the photo of the Japanese trade minister, and a screen shot from C4 news with a black presenter. Still though, the imagery is not a balanced one. For instance, there are very few (positive) images of women through the magazine. I would like to make a suggestion that you could for instance have a column from linuxchix or similar groups and try to foster a more mixed content.

If Linux is to become the world's desktop, the majority of users will not be white. In fact in the near future this would probably come about due to mass adoption in India and China.

Open Source is for everyone. Britain is a multicultural nation. The images in *LXF* should reflect this.

★ Letter of the month

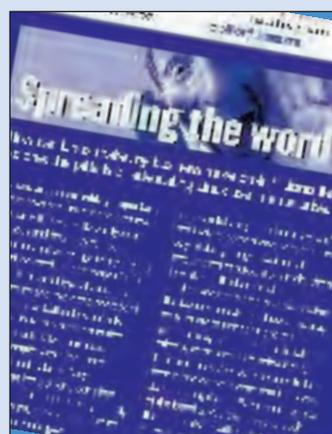
This month's winner receives a copy of *The Cathedral and the Bazaar* by Eric Raymond (ISBN 0-5960-0108-8)

Linux needs marketing

I work in the IT Department of a retail company and since I made my first steps in Linux, I knew that Linux and the Open Source software was the right way to get a great advantage over our competitors. But the company (Directors, President, etc) are not interested in a change or even probe the new approach and surprisingly re-enforce the relation with Microsoft by migrating to Windows XP.

Each store performs all its operations through applications in Informix DBMS running on Unix Servers (clusters or dedicated servers). Obviously they already know about the restrictions of another OS.

Are they so rich? No. The company cannot change to an entire ERP system (like SAP, JDedwards, etc) because of the



Effective advocacy of Linux is just as important as coding applications!

high costs. The company still has to produce its own software (payroll, retail apps, prices, etc). Obviously using Open Source software would be a great solution to make our business better.

I live in a third world country, and like Jono Bacon commented the other month, Linux and the Open Source software is the kind of

technology we need to reduce the digital gap, but sadly the media and advertising never show news or products from Linux or Open Source software. Amazing changes like the German or Brazilian government will never get enough noise in the media, but advertising, deals and promotions with Microsoft always get the first place.

So, please advise me and other readers in my position about a good book, site or tactics to improve the Linux marketing or guide me in the creation of a LUG inside the company.

Juan Panexxo

Well, as you mention Jono, I suggest you also check out his advocacy advice that appears on his site and forums at www.jonobacon.org. As for books, the best work for understanding Open Source is the one we're sending you! For more about starting a LUG, www.linux.org/groups/ is a good place to start, and there's some info at www.lug.org.uk/ may be of help to you too. Good luck!



Linux should not be just the white man's OS.

Mel Glynn, via email

The people pictured, for example in the news section or in relation to interviews etc, are obviously reflective of the readership. Where we occasionally use stock imagery as feature openers and page furniture in the magazine, the main criterion we judge on is the suitability of a picture in

In a 99% male dominated market, is it surprising we don't get many pictures of women in the magazine?

relation to the editorial it accompanies – race or gender of the image subject doesn't enter the equation at all. We have included women and, as far as I can remember, people of various races – to be honest, it has never occurred to us to deliberately include someone's face just because of their race, and we don't keep a running tally.

I'm not sure that featuring more women in the magazine just because they are women is necessarily the right thing to do either. We welcome contributors to the mag whatever their gender, so if you can write about Linux, whoever you are, get in touch!

LINUX ON LAPTOPS

Why vendors are rarer than hens' teeth...



Transtec offers a range of Asus laptops preloaded with Linux – see its website for more details.

Laptop as console

Having worked with Unix for many years, I'm heartened to see Linux entering the mainstream. However there's still too many hardware/ components, modems, printers *etc.*, out there which are designed for 'non-operating systems' like Windows XP, Windows 2K *etc.*

I've a TARGA laptop onto which I managed to load Red Hat 9 and with a GRUB loader was able to keep WinXP and Win2K partitions, for a couple of apps that I haven't been able to find Linux versions of yet.

Anyway, there's a few bits of hardware in it, like the IR port and power management, that I wasn't able to get working. Also the machine is 18 months old and in the last 4 months it's been back to Germany three times for repair. First, it went totally dead and had to have a new motherboard fitted, then when it came back, the CMOS wouldn't hold settings, even with a new battery; so another motherboard was fitted but even then XP and 2K would simply hang after a few minutes work. Glad to say though, Linux would run perfectly!

Getting to the point – I'm in the market for a new laptop therefore and would welcome LXF's or its readers comments/recommendations as to which brands of laptops are the most Linux-friendly.

I'd like a laptop with no Winmodem, hardware that Linux can probe and which don't insist on pre-loading the full disk as a single partition full of Microsoft's bloatware.

I'd also like a big hard disk – greater than 50 GB, DVD burner, WiFi and Bluetooth; and I don't want to have to buy USB convertors just to be able to plug a serial cable into a larger Aix/Unix systems to act as a console.

John McQue, *via email*

Pre-loaded Linux

Now that Evesham is not offering a Linux system as part of its product range, would you please advise me where and from whom I might buy a laptop or desktop in this country.

Peter McCourt, *via email*

We've had several emails in this vein over the last few weeks. Sadly, there is no UK High Street vendor currently offering a laptop pre-installed with a Linux OS of any description directly to the public. Business users fare little better – some business-oriented hardware vendors will pre-install the Linux OS of your choice on quantities of portable PCs ordered from them. But, even assuming they have a reasonable level of Linux know-how, in a lot of cases a vendor will refuse, as despite anti-trust rulings in the US, it appears that Microsoft still suggests to European vendors that cheap Windows licenses are only available to retailers that sell hardware exclusively equipped with it. This state of affairs might be why some of our readers have had problems in the past obtaining a laptop without *any* OS pre-loaded on it.

The UK arm of German company Transtec sells a range of Asus Laptops preinstalled with a choice of Windows XP, SUSE 8.2 or no OS, as shown above. For more, see www.transtec.co.uk and www.transtec.de. We hope to bring you a review of some Transtec hardware in a forthcoming issue. If any readers hear of any other UK vendors offering Linux laptop or desktop PCs for sale, please do drop us a line to us at the usual address so we can tell all our readers!

It's quite straightforward to install the major Linux distros on modern laptops – see our Megatest in *LXF41* for a comparison of six of the best, and get more info about installing on many portable PCs at www.linux-laptop.net/.

NEW: SUSE LINUX 9.0 – the complete suite for newcomers and professionals. Also runs parallel to Windows!

SUSE LINUX 9.0: An Operating System with 1,000 possibilities, not a 1,000 vulnerabilities

With SUSE LINUX 9.0 you get a complete Desktop Operating System with all the productivity and multimedia applications you need. In addition, its secure system architecture and built-in firewall will free you from the annoyances you've come to expect from the Windows world of crashes, viruses, and worms. Also available for the first time in 64bit version for AMD64 processors. Get the facts now.

Also available at Amazon.co.uk, dabs.com, PC World

WWW.SUSE.CO.UK

« UK Linux PCs?

After reading LXF43 I became intrigued by the Evesham Micro desktop running the Debian OS that you reviewed with astonishing ratings in the magazine. I am a Dane and I sent a relative out to buy an example for me at the beginning of October and guess what the result was: "Evesham had pulled the desktop from their product list after two months (in early September) because of low sales and they had no plans to sell a Linux-based desktop package in the near future".

SUBMISSION ADVICE

WHAT WE WANT:

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

WHAT WE DON'T WANT:

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

WRITE TO US AT:

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



If any UK retailer wants to fill the gap in the Linux hardware market left by Evesham's exit – now's the time!

Oddly enough, they will sell me a similar specification computer with Windows XP at the approximately same price but no printer, no speakers and the flat screen terminal is no longer flat. It seems to me that the task of delivering the desktop with the Debian OS or with no OS installed should not be a major task. It must be frustrating to review a desktop giving it top rating, a rating I don't remember having seen before in the magazine and having it taken off the product list a month later

Ove Winther Jensen

As in the laptop market, business users have it a whole lot easier when it comes to buying machines with Linux pre-installed. In 2001, Dell used to offer

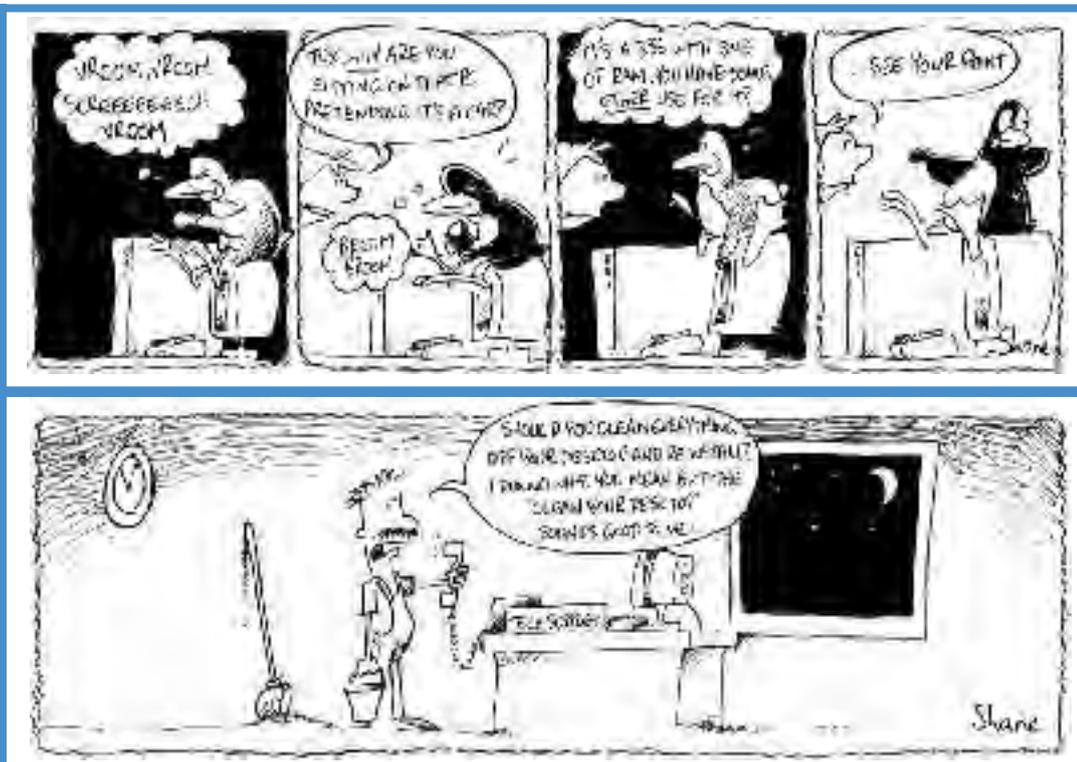
desktop PCs with the option of Red Hat, but this was discontinued. Microsoft execs internally touted this move as a triumph for MS anti-Linux strategy, according to emails later revealed by Microsoft's antitrust trial. Dell still offers Linux on some workstations, but on desktops and laptops it is only available pre-installed in batches of 50 or more. In July 2003, HP introduced the HP Compaq Business Desktop d220 Microtower, aimed at small and medium-sized businesses, with either Windows or Mandrake pre-installed. HP also offers other desktops pre-installed with Linux, including the recently announced HP Compaq d530.

When we first contacted Evesham on our readers' behalf, the company still

maintained that it was committed to a range of Linux desktops – some time after we reviewed the products. But as we were going to press in November, we received an email from Evesham stating: "The sales of the Linux systems have been disappointing in this quarter and we are temporarily withdrawing the Linux option." Should any prospective Linux desktop buyers want to show Evesham that there is a demand for its Linux desktop computers by emailing customer.care@evesham.com, if you CC lxf@futurenet.co.uk with any mails that you send, we will be able to monitor the demand situation more easily. Alternatively, you can speak to someone in Evesham's customer care department on +44 (0) 870 160 9500

Helpdex

shane_collinge@yahoo.com



READER TIPS

64-BIT IN THE HOME!

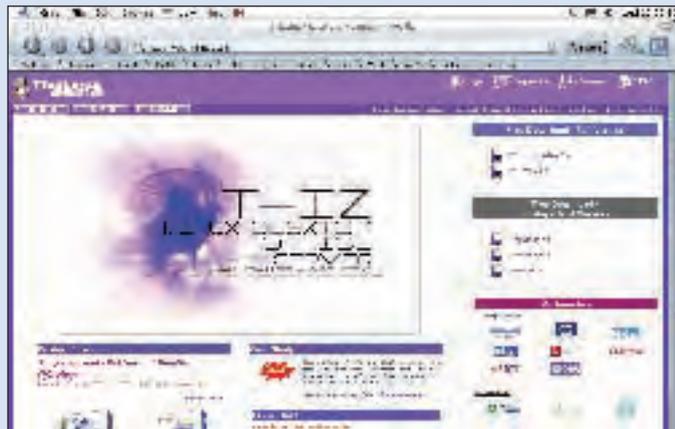
Many mags have been covering the release of AMD 64-bit processors recently, and there is no doubting that they are very impressive. For anyone that is dreaming of such a beast, did you know it is possible to buy a 64-bit computer on eBay for less money than a simple PC? It has good performance, running a solid Unix made by a famous manufacturer – Sun. The Ultra series of machines have now been replaced by newer models, and these are now widely available second-hand. Ultra 5s and 10s use PCI expansion cards and IDE drives, so they should be familiar to anyone used to modifying a PC. Sun has a very healthy attitude to software licensing and there are Solaris news groups, so support for home users really is pretty good.

I have just ‘taken the plunge’ with an Ultra 10 and found it very easy to install Solaris and configure the machine. As an alternative to making a standalone Linux box it really is worth some serious consideration (I have tried both routes), and although faster machines are available today, these are not so very much slower and certainly capable of anything you may wish to do at home. Take a look, you may be surprised – I was! Phil Macphail, *via email*

DIGITAL SPAGHETTI

Just thought that I should comment on the answer to one of the letters in *LXF45*, as you seem to have fallen foul of a popular misconception.

You are not alone, as I have seen the same error in textbooks and elsewhere on the Net, etc. According to Gilbert Held’s book, *Ethernet Networks* (ISBN 0-471-25310-3), and several other reliable Web sources, such as: www.darron.net/network/fifthpage.html, www.intel.com/support/express/hubs/100base/21370.htm, www.dewassoc.com/support/networking/buildfastethernet.htm



The 3-4-5 rule only applies to 10Base2 and 10BaseT networks. In 100BaseTX networks, the maximum for classII hubs is 2 per network, and for classI hubs (not common, most hubs are classII) is 1.

Also, the maximum cable between 2 classII hubs is 5 metres, limiting the total network span to 205 metres. The only way around this is to stack the Hubs, or use a Switch, and nowadays Switches seem to be almost the same price as Hubs, or even cheaper, I bought an 8 port switch last week for #14.50 +VAT! (yes, I know it won’t be top-quality at that price, but it was for a two-person workgroup with a networked printer).

The other consideration is cabling. Stranded core ethernet Cat5 UTP is only recommended for cable runs of less than 10 metres, due to increased resistance and crosstalk. Over 10 metres, solid core cable is recommended.

I install MFPs on networks on a regular basis, and have built and administrate a few small (less than 30 users) office networks. In my experience, I have found that the most common causes of 100BaseTX problems are:

- 1 Incorrectly wired cables (eg using pins 1+2, 3+4, 5+6, 7+8 instead of 1+2, 3+6, 4+5, 7+8)
- 2 Badly wired and/or untested cables
- 3 Too many Hubs/Switches or badly laid out networks (switches off hubs off switches etc.)
- 4 Overlong crossover cables between two machines.
- 5 Using stranded cable for long runs.

Number 1 is quite important, as 100BaseTX uses wires 3+6 for collision detection as well as receiving data, so though it would work with 10baseT half duplex, it would not work with 100BaseTX. Leo Maxwell, *aka WylieCoyoteUK on LXF forums*

THIZ LINUX

I got a rather Odd Distro today, I’m using it now. It comes with a Chaintech (I think) motherboard as a freebie.. Its called ‘Thiz Linux’ – upon booting it Looked like my old

Visit www.thizlinux.com and www.chaintech.thizlinux.com for more ThizLinux information.

faithful RH7 Distro, but its a Hong Kong-based distro – very weird, 2.4 kernel, office app, CD-R Burn app, KDE, lots of other goodies.

So I booted (I had spare 5GB partition) and almost without asking, it installed itself, no partitioning, formatting, building or anything . I was worried it was going to take over the whole HD, but on rebooting i was pleasantly surprised as it (correctly) installed a pretty bootloader, select Windows or Linux – it even uses the ReiserFS. Yum yum! It would seem to me to be the best distro for newbies ever that I have seen, and I’ve been around since Red Hat 4.3. It may make a superb ‘get you started in Linux’ disc – no config, boot straight into X – far less stress than that other OSES to install. The only thing that went wrong was that it didn’t find the SBlive 5.1 digital first time... Well, I suppose you can’t have everything! Love the mag, let’s have more config stuff!

Mark, *Morecambe, Lancs* [LXF](mailto:mark@morecambe.lancs)

GRAPHICS CARD NOT HARD!

One problem, two solutions

Re: Hans’s problem *Graphics Card Hard* that was in your *Answers* section, *LXF46*, page 97.

Your problem is pretty much the same as what I experienced recently, and is due to the lack of AGP3 support in the 2.4 kernel. When an AGP3 card is detected, the chipset switches to APG3 mode, causing problems for the 2.4 kernel.

Currently, it would seem you only have two choices, either use the 2.6 test kernel, or assuming your motherboard chipset is the VIA KT400/600, get the 2.4.21 kernel source and apply the patch available here:

<http://lists.insecure.org/lists/linux-kernel/2003/Mar/att-3999/agp3-2.4.21-pre5.diff.bz2>
Shane, *via email*

Reply to *Graphics Card Hard* (*LXF46*, page 97): the problem with Mandrake 9.1 on the computer in question is not the graphics card, although that would usually seem to be the most likely culprit.

I noticed in the config he sent in, that he has 1GB of memory. I had the same problem after upgrading my memory to 1GB. The Mandrake 9.1 workstation installation does not have large memory support set up in the kernel, and fails on a machine with more than 512MB of RAM. I solved it by removing 1 stick of 512MB, configuring the Kernel for large memory support, and then reinstalling the memory stick.

Leo Maxwell, *aka WylieCoyoteUK from the LXF forums*

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/Documentation: Whichever is most appropriate!

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.



4-5 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...

Slackware 9.1

The latest release of Linux's very first distribution maintains its stable reputation **p20**

HP Integrity server

The first Itanium server we've managed to get – how does it compare to Opteron? **p22**

Gigabyte server

The SR125ED improves upon its predecessor, the well-regarded SR113 **p24**

GAME: Majesty >>

Commercial games ported to Linux are few and far between: this first RTS release from Linux Game Publishing hopefully could set a new trend? **p26**

Books

Subjects include C++, intrusion detection and TCL/ TK this month; plus a battle between two Red Hat behemoths **p28**



LINUX FORMAT BENCHMARKS EXPLAINED

To provide objective performance comparison between machines running Linux, we run a set of in-house benchmarks. These are: *bonnie* and *hdparm* to test hard drive performance ('hd' in the benchmarks), MySQL *Super-Smack* to test how well a machine handles database serving ('mysql'), *ApacheBench* to test how fast a machine can serve web pages ('apache'), a *gcc* compilation of Linux kernel 2.4.19 ("compile"), and *oggenc* to convert a test .wav file to a .ogg file. These numbers are then averaged to produce an overall score, which may be adjusted

slightly now and then, if a machine has a particular high or low point that should be taken in to consideration. Combined, these tests really push hard drives, network cards, and CPUs to their limits, and so give quite a representative figure – a multiple of the performance our yardstick machine.

The *LXF* yardstick box attempts to represent an 'average' reader's setup: Debian 3.0 on an 866MHz PIII with 256MB of RAM. So, a machine which scores 1.5 on our *Apache* test served 50% more web pages than our yardstick, whereas a box that scores 3.0 for overall ran, on average, three times faster than our yardstick box.

BENCHMARKS

HD	0.83
APACHE	1.22
MYSQL	1.11
COMPILE	0.96
OGGENC	1.71
OVERALL	1.17

The blue bar in the example above represents the performance figure for the hardware, and the red bar is the benchmark figure. When a piece of kit performs lower than the benchmark, as in 'hd' and 'compile' above, the blue value will appear less than the red value. **LXF**

X86 LINUX DISTRIBUTION

Slackware 9.1

Mike Saunders puts aside his Cumbrian ale and harpsichord music to check out another venerable classic: Slackware.

**BUYER INFO**

Vector Linux (also freely available) is based on Slackware, and aims to improve the ease-of-use issues.

- **DEVELOPER** Patrick Volkerding
- **PRICE** Free, or \$39.95 for box set
- **WEB** www.slackware.com

There's an old saying in the Slackware community: "If you use Red Hat you'll learn Red Hat, and if you use SUSE you'll learn SUSE, but if you use Slackware you'll learn *Linux*." While it's true that all of the major distros can be pulled apart and examined, Slack has a long-running reputation for being cleanly laid-out, well designed and superb for learning the workings of a Linux system.

It's also one of the most mature distros – the earliest releases appeared in 1993, and a decade later, lead developer Patrick Volkerding is still doing the vast majority of the work. This one-man operation means that the software package count doesn't always rival the bigger distros; the upshot though is a cohesive design, top-notch stability and a highly understandable system.

Slackware 9.1 is freely available for download, and this release has jumped up to two binary package CDs (the second containing GNOME and KDE). Alternatively, a snazzy 4-disc boxed set can be bought from <http://store.slackware.com> – this includes the source archives and ZipSlack mini-distro, also downloadable, and it's a good way to support the project financially.



A vanilla KDE setup is one of Slackware's available WM/desktop choices.

With its no-nonsense text-based installer, Slack's system requirements are refreshingly low, and a 486 with 16MB RAM should suffice. The menu-driven setup process steps through partitioning (courtesy of *cfdisk*), package selection and network configuration, and getting up and running on a modern box takes less than half an hour. Newcomers may find the lack of setup wizards and automagic hardware detection off-putting; still, the well-commented config files and addition of the 'hotplug' subsystem helps greatly here.

Slackware 9.1 sports an impressive range of the latest stable goodies:

GNOME 2.4.0, KDE 3.1.4 and XFree86 4.3.0 sit atop a rock-solid base of kernel 2.4.22, *glibc* 2.3.2 and GCC 3.2.3. On the server side, *Apache* 1.3.28, *OpenSSH* 3.7p2, *Sendmail* 8.12.10 and *MySQL* 4.0.15a are included, while developers will find Perl 5.8.0, Python 2.3.1 and Java 2 SDK 1.4.2. A healthy range of up-to-date packages, then, and although it's not as broad as Debian and co, various sites offer a range of extra add-ons.

Slack to the future

Compared to the relative complexity of RPMs and .debs, Slack's gzipped tarball packages can appear old-hat at first glance. They're nothing more than .tgz archives with some extra info (description, post-install scripts etc). Yet this is exactly what's intended, and what Slackware users love – simplicity. It's a breeze to manipulate, change and create packages, and the new additions of *Swaret* and *SlackPKG* make dependency resolving and clean upgrading *à la* Debian very straightforward. Pleasing to see.

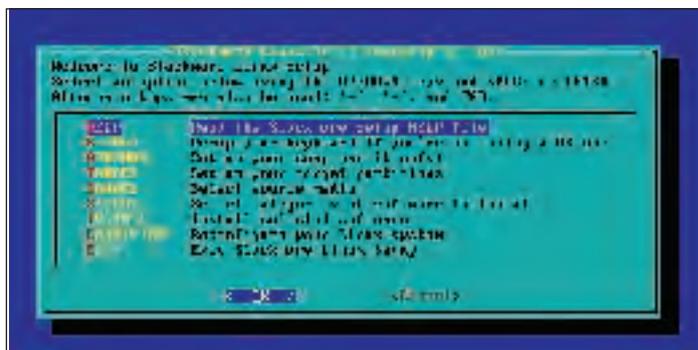
In day-to-day use, Slack 9.1 is a spritely performer thanks to its 686 CPU optimisations (but 486s and

Pentiums etc will also work) and bootup times are about half that of a stock Red Hat 9 installation. The sane and clean filesystem structure and BSD-style startup scripts help; Volkerding's policy is to leave software installation as the original coders intended, and this certainly makes life easier when installing or upgrading bits by hand from the source.

As a result of its consistency and reliability, Slackware is commonly used in production servers, and so security fixes are a critical issue. Volkerding aims to support the current stable release along with the one previous, giving each release approx 12 months of fixes (and sometimes more).

On the whole, Slackware differs from most of its brethren in one fundamental way. While Red Hat, Mandrake and SuSE *et al* all strive to deliver a 'Linux experience', which is clearly a bonus for first-timers, Slack essentially says 'Here's some solid and tested software, all packaged up and guaranteed to work together, so go make the Linux system you want'. It stays out of the way, providing a very stable and easily understandable base to build from.

And consequently it's difficult to give Slack 9.1 a definite score; experienced users will love the speed, simplicity and reliability of the distro, and so it deserves a 9/10 in this case; but equally, some newcomers may find it too demanding and a little unusual in places. So we think 8/10 is appropriate, and strongly recommend anyone with moderate Linux experience to give it a try. **LXF**



Slack's lightweight and workable installer gets the job done without hassle.

LINUX FORMAT VERDICT

FEATURES	8/10
PERFORMANCE	10/10
EASE OF USE	6/10
DOCUMENTATION	8/10

Stable, simple and secure – Slack's mantra. Another rock-solid and speedy release, with plenty of the latest goodies, but novices will need the FAQs and HOWTOs...

RATING **8/10**





HIGH-END SERVER

HP Integrity rx2600

Already on its second generation and quickly approaching the third, Paul Hudson finally got hold of an Itanium, and boy, is he impressed...

BUYER INFO

A high-performance solution for the most demanding environments. Those with less cash to hand might be tempted to consider a Xeon- or Opteron-based solution instead.

- **SUPPLIER** Hewlett Packard
- **PRICE** £22,930+VAT
- **WEB** www.hp.com

Years before Jerry Sanders had his plans for AMD64, Intel's own 64-bit plan was already brewing. The Itanium family of CPUs was a drastic change for Intel, as it was built the expertise HP gained while producing its PA-RISC line. Technologically speaking, it has little in common with the x86 architecture most of us are familiar with, and instead is a complete redesign from the ground up, designed for maximum performance in a modern computing environment.

HP – being the co-designer of the CPU – does an unsurprisingly good business making and selling Itanium machines that are usually geared towards high-performance computing environments. This machine is a 2U rack-mountable system that comes

packed with two of the new Itanium 2 CPUs running at 1.5GHz apiece, both of which have 6MB of L3 cache. On board is HP's exclusive zx1 chipset, which was designed to eke every last drop of performance out of the Itanium chips. Backing all this up is 12GB of DDR RAM, two 36GB Ultra320 hard disk drives, and a DVD-ROM drive.

Into Itanium

This is the first Itanium machine we've reviewed here at *LXF*, so before we take a closer look at its performance we wanted to briefly explain the architecture of the machine. The key difference with Itanium is that it follows the EPIC architecture, which stands for *Explicitly Parallel Instruction Computing*. This paradigm is designed to eliminate the mistakes seen in x86, where the CPU did peephole optimisation of code before it was executed – it literally examined a very small set of instructions, and re-ordered them on-the-fly for maximum performance.

As you can probably imagine, this optimisation takes up quite a bit of die space on the CPU and also isn't entirely conducive to maximum performance. With Itanium, this no

longer happens – if you use a fairly basic compiler, your code will run at less than 33% of its potential, which often leads people to claim that the Itanium is slow. However, if you use a compiler such as *icc*, Intel's own compiler, you'll almost certainly see your performance shoot through the roof because the Itanium architecture is very well designed.

The one drawback to Itanium has always been its x86 compatibility – there are very few companies willing to leave behind all their x86 software and make the jump directly to Itanium. Sure, it helps a lot that our Linux distros come with source code, so we aren't affected by the majority of porting issues, but there are a substantial number of enterprise apps, notably Oracle and other databases, that need special Itanium versions. To counter this, there is x86 emulation inside Itanium which runs at notoriously slow speeds, despite much work being done to try to improve it.

So, all in all the Itanium is a machine that *can* run x86 code, but would rather not – the question is, when running 64-bit code, does it actually perform as well as everyone seems to say it does?

No LXF benchmarks

Itanium machines are, to be blunt, far too expensive to be put to work as a simple web server or Ogg encoder. Having said that, we still do endeavour to perform as many of our benchmarks as possible in order to at least give you a rough idea of the relative performance.

Sadly, the machine we reviewed comes bundled with Red Hat Enterprise Linux AS 2.1, a product containing lots of fairly out-of-date software that has been recommissioned for use on the Itanium. While this does make for a lot of stability, it also makes it nigh-on impossible for us to actually run our standard benchmark suite on the system, simply because the software supplied is either really old or poorly configured. For example, we were somewhat surprised to learn that it comes with *oggenc 0.8RC2* installed by default, which even novices will agree isn't really a version number that inspires confidence; and, lo and behold, it failed to compress our file.

Although we can't print any real benchmarks, we can at least vouch for the fact that the server is very fast for general use – not a very objective measure, perhaps, but all we can offer.



The solution here is to ask for SUSE Linux Enterprise Server in place of Red Hat's offering, as we've tested SLES extensively in the past and have found it a great deal easier to use and doesn't insist on installing ancient software versions across the board.

Much can be drawn from the sheer statistics of this machine – two U320 SCSI hard drives coupled with 12GB of RAM and the very latest Itanium 2 chips will undoubtedly provide performance to be proud of on all fronts. Remember, each of these Itanium chips has 6MB of L3 cache, which alone makes them between 20–50% faster than the older Itanium chips. Of course, all this quality adds to the cost, but you get what you pay for.

Mighty miscellany

Too many review machines we see here at *LXF* come in a box devoid of any documentation or even CDs. So, it was a refreshing change to see this HP server arrive with a selection of manuals, CDs and cables to help us get set up. The

Quick Setup guidesheet and Getting Started Guide booklet complement each other nicely, and provide a short introduction to the machine and its capabilities. The Getting Started Guide is particularly good, providing twenty-two pages of troubleshooting as well as other helpful installation tips.

Most impressive of all, though, was the *HP Enablement kit for Linux* for Itanium 2-based Servers, which is a CD ROM and booklet that helps you plan and prepare an installation of Linux (Debian, no less!) on your Itanium box. This is backed up by another two HP CDs, *Itanium Processor Family offline diagnostics and utilities*, which is designed to boot directly from EFI in order to solve otherwise fatal problems, and *Documentation library CD*, which contains much of the paper docs in electronic format, but also a selection of other helpful guides and walkthroughs. In addition to all the above, you also get a full copy of Red Hat Enterprise Linux AS 2.1, with manual and all five CDs. On the hardware front, you get a fairly standard mouse and keyboard combo, as well as a variety of other cables 'just in case'.

Aesthetics

As far as documentation, CDs, and accompanying hardware goes, this is as good as you'll get anywhere you look, and reflects very well on HP's quality of service. It's a shame really that so few other companies even try to reach this same level of thinking – we've looked into too many boxes for £5000+ machines and been sorely disappointed to see nothing but a large chassis.

Because the Itanium is a wholly new architecture as opposed to being built on any ancestral requirements, it's pretty much 'cruft-free' in that everything on the chip die is there because it needs to be. As a result,

Itanium is very well designed to be economic when it comes to space, and also therefore when it comes to heat. Despite having a total computing power of 3.0GHz – and, please, try not to equate that to the performance of a 3.0GHz Pentium 4, because they are in entirely different leagues – the HP server is remarkably quiet, even when under heavy loads.

Another stereotype happily smashed by HP is that rack-mounted servers need to look bad. Although this machine certainly wouldn't be welcome on a desktop, it has at least had some degree of effort expended upon it to make it easy to administer – from the small writing on each hard drive informing you of its size speed and type, to the pull-out support information card – it all combines to make the machine that little bit more user-focused.

Around the back there's two 10/100 NICs and a gigabit connection as well, plus a redundant power supply, three USB ports, and the usual array of connections (VGA, etc). On the front and back is a hardware locator light, which is something you tend to only see in luxury servers such as this one. If the concept is new to you, have a think how you'd locate one particular server amongst hundreds – wouldn't you like to be able to send a command to that server to start flashing a very bright, blue LED to help you find it?

Shamed by software

Without solid benchmark figures, it's hard to tell quite how well this server performs in comparison to our other machines, which is an annoyance and a worry unless you've already committed to the Itanium platform. As servers go this is certainly a thing of beauty – nothing has missed the attention of the designers, and this is very encouraging when it comes to

considering the quality of the components inside.

When it comes to the software pre-installed on the machine, it's a shame it's so hard to use – our benchmarks aren't complicated by any means, yet we still couldn't get them working properly. In our experience, SUSE Linux Enterprise Server wholly outclasses Red Hat in terms of software configuration and ease of use, and you'd think HP would have put two and two together and signed some deals already.

With what there is currently, this is a great machine that could be a lot more if more configuration was done before release. HP is well-known for producing the finest systems on the market and this is clearly shown in the finesse built into this machine. The price tag is high – very high – so you really need to be committed if you want to purchase a rack of these things. It also shows that *LXF* could rework its benchmarks so that they can work on older software if we want to be able to accurately check the performance of this beast.

So, in conclusion, this is a machine that exudes design excellence and advanced technology, but is let down by frustrating software and sky-high price tag. Is it worth the money? Yeeees... but only just. **LXF**

LINUX FORMAT VERDICT

FEATURES	10/10
PERFORMANCE	9/10
EASE OF USE	7/10
VALUE FOR MONEY	7/10

Itanium is a great architecture, but this machine seriously lacking full support with the bundled software.

RATING **8/10**



Just to the left of the collection of USB ports is the blue hardware locator LED. You are of course advised to remove the key from its housing on the right of the backpanel prior to deploying this server in any environment where physical security is an issue.

1U SERVER

Gigabyte SR125ED

Server supremo Gigabyte is back with its latest and greatest 1U system.

Paul Hudson puts it to the test...

BUYER INFO

An excellent, if a little top-heavy 1U server that's very aggressively priced. If you're really cramped on space, consider Opteron-based machines such as the Armari 1U server reviewed in LXF45.

- **SUPPLIER** Upgrade Distribution
- **PRICE** £1488+VAT/£1797+VAT
- **WEB** www.upgrade-distribution.com
- **PHONE** 01252 332800

The last Gigabyte server we reviewed was the SR113 back in LXF40, and we gave it 9/10 and a Top Stuff award. Quite rightly, too - the machine was well-designed, feature packed, and still managed to come in under £1000, which given the high specs of the system at the time (single 2.4GHz P4, 1GB RAM, and 3x30GB hard disks) was, and still is, a fantastic bargain.

Although the SR113 fills a need

in the marketplace, what do you get if you take the same extreme attention to detail and the same aggressive pricing that Gigabyte is renowned for and ask the company to produce a model suitable for more demanding applications? The result is the SR125ED, a machine that comes backed with two 2.8GHz Xeons, 1GB of DDR RAM, and four ATA-133 80GB hard drives. This is quite a step up from the £872 system we reviewed back in LXF40, and so the price tag is correspondingly higher - £1797 + VAT.

Panoramic view

The SR113 machine we reviewed back in LXF40 was a real cracker on the performance front, particularly when you took into account its incredibly low price tag. Switching to a dual-CPU system here accounts for most of the price-jump, but the added hard drive

and all-round increased capacity (320GB when used with plain RAID striping) is also most welcome.

The machine, provided for LXF by long-time Gigabyte experts Upgrade Distribution, didn't come installed with any version of Linux, or indeed any OS at all. We had to put a distro on ourselves, which caused a minor hitch because our distros didn't actually have support for the RAID card used in the machine. Fortunately, Upgrade Distribution continue to show their mastery of server design and bundle with this server a USB floppy drive and Linux driver disk that make installation a cinch.

The build of this machine is largely similar to the SR113 - there are quite a few USB ports around the front and back, plus CD-ROM drive and two network interfaces. Although we usually avoid making any changes to

server machines we get so as to avoid any claims of bias, one minor tweak we did make to this box was to enable hyperthreading. HT technology has been shown to be as much of a hindrance as a help across balanced tests, so, to ensure this machine is equally helped and equally hindered when compared to the other machines that we benchmark, hyperthreading was enabled. We recommend you run your own tests to evaluate whether hyperthreading is beneficial in your environment.

Sunshine in a bag

With more than twice the CPU performance compared to its younger cousin the SR113 as well as beefed-up hard disks, we anticipated that this machine would easily be able to turn out an above average performance result. Looking at the numbers, there are a few highs and a few lows, but the general shape of the results corresponds to what we'd expect from this class of server. That is, this server has a lot of CPU power and RAM, but uses IDE drives to cut back on cost -

Less connectivity than some other servers we have reviewed, but when you consider how this hardware will be deployed, this isn't really a problem.



no bad thing in itself, but it does make the machine a little top-heavy. This is minimised by having four hard drives that you can chain together as a stripe set in order to maximise performance, but at the end of the day this machine would really benefit from SCSI.

The biggest 'low' is in the PHP/MySQL test, which returned a result of 0.75 – quite disappointing, really, given that there's a lot of CPU power backing this thing. The problem here is easily put down to the hard drive performance again, because the PHP/MySQL test does a lot of random seeking – something at which SCSI drives excel. You can try and improve this somewhat by tweaking the stripe size on your drives to fit the kind of task you've assigned to the server.

Both the *oggenc* and *Apache* scores are particularly good, as can be seen simply looking at the figures, but these are dragged back by the score for *kernel* compilation – this score is usually very telling when it comes to real-world tests because compilation is an arduous task that stresses a variety of CPU functions. In this situation, a

score of 1.59 is a little below average for a machine of this spec, and hints that you really need to tweak the settings yourself to get the most for your money.

Overall, this is certainly a machine suited to particular tasks, and is likely to be well at home wherever hard disk performance isn't key. Clearly, that precludes this machine from the role of file server if you are looking for a departmental machine, however rack-based servers such as this one are generally used in HPC environments where CPU, RAM, and network performance are key to excellence. Fortunately for this machine, that's where it excels – this server is exceptionally fast at raw processing, which makes it a particularly good choice if you want to cram in as much performance as you can in a small space. At 1U, even a half-rack of these will scream thanks to the build quality and inherently fast components.

Pick and choose

The key benefit of purchasing a server from Upgrade Distribution is that you can choose each component that goes into your server right down to the

smallest level – if you want fewer drives, or even a one-CPU SR125ED the manufacturer can build the exact machine for you. The SR125ED single-CPU version starts at £1488, however you'd be well advised to go for the two-CPU option reviewed here at £1797. This is further helped because Upgrade like *LXF* readers so much (hey, let us dream...!), you can get a further £50 off this price if you quote LINUX03 which makes the price even more of a bargain.

Having said that, you may find this server a little noisy, particularly if you have a rack of them running in a room with little noise shielding. The reason for the noise is because Gigabyte hasn't held back any expense when it comes to cooling – there are four high-power fans onboard that keep the whole box running icy cold.

As with the SR113 we reviewed, this machine came with little or no documentation – not a crushing flaw by any means, but a significant irritation nonetheless. The hard drives were packaged separately from the computer and not stickered '1', '2', '3', and '4', which pretty much voided any possibility that a distro had been pre-installed for us. Still, it's true that the drives are safer when packaged separately, so you can't hold it against the manufacturer or supplier.

Although the USB floppy drive is a neat way of solving the problem, the server also comes with a converter kit to allow you to fit your own floppy drive in place of one of the hard drives. With floppy drives on the way out, this is probably overkill, but it's nice of Gigabyte to give its customers that extra choice.

Feeling glad

This is definitely a worthy successor to the SR113, and provides a logical step up in terms of performance. The little touches such as the inclusion of the USB floppy drive are the hallmark of Upgrade Distribution's quality process and are what, all considered, make this machine a safe purchase. Yes, you can get a lot of the hardware yourself off the shelf, and if you haggle hard you might be able to get roughly the same price, but this machine comes with the little extras that ooze forethought from Gigabyte, which is very reassuring.

On the downside, the hard disk performance lags a little because of the choice of ATA drives. This is an inevitable choice for a low-priced server such as this one, but there's

hope yet – Gigabyte are currently evaluating a move to S-ATA sometime in 2004, which should help things along a little, particularly as the S-ATA standard moves on.

To improve this server, Gigabyte really needs to get into the area of documentation writing. As it stands, you get this machine, its drives, and the USB floppy all in the box, but next to nothing when it comes to how to put it all together and get it working. Yes, veteran sysadmins will have no qualms about assembling the machine themselves, but it really would only add a little cost to the system for Gigabyte to print some helpful documentation and troubleshooting guides. The simple fact is that when you pay £2000 for a server, you expect at least a little paperwork to accompany it!

Other than the complete lack of documentation, the only thing holding this back is the drive configuration, and this is a simple equation of 'you get what you pay for'. From the high-performance computing perspective, this is exactly what the market needs to keep up the pace of processing. Having said that, although the main processing setup here is so powerful that it can go a long way to masking the hard disk performance, if you really, really need that extra oomph from your hard drives, then this server will disappoint. Otherwise, this is another sterling machine from Gigabyte that keeps true to its long line of successful products and is sure to delight long-time Gigabyte fans. [LXF](#)

BENCHMARKS

HD	1.62
APACHE	3.25
MYSQL	0.75
COMPILE	1.59
OGGENC	3.87
OVERALL	2.22

LINUX FORMAT VERDICT

FEATURES	9/10
PERFORMANCE	8/10
EASE OF USE	10/10
VALUE FOR MONEY	8/10

Packs a powerful punch for such a low-cost system, and is sure to keep customers happy for years to come.

RATING **9/10**



GAME

Majesty

It's not just the cream that floats to the top... Ascending to his rightful position, **Nick Veitch** sets about his own reign of terror.



BUYER INFO

Real-time strategy game ported from Windows. Also try Kohan for more traditional Linux RTS thrills.

- **DEVELOPERS** Cyberlore
- **PORTED BY** Linux Game Publishing
- **PRICE** £27.00
- **WEB** www.linuxgamepublishing.com/

Commercial Linux games are quite rare these days, but not because of any shortage of people to play them. Did you know that gaming is altogether the number one reason for people dual-booting rather than just ditching Windows? But there are some great Linux games, and *Majesty* hopes to extend that family by at least one.

The obligatory plot surrounds a new king who must restore the realm to its historical position of power and peacefulness. Against these lofty aims are ranged the ambitions of evil wizards, the bestial nature of minotaurs, the arcane powers of the undead and various other things too unpleasant to mention, some of whom fashionably sport their entrails on the outside of their bodies.

One of the key differences between this game and a slew of other fantasy RTS games is that you



The wizard's get some shock and awe from a passing dragon, but magical countermeasures are at the ready.

don't actually control any units. You can hire 'heroes' (see *Holding on for a hero* box, below), but they decide for themselves where their interests are best served. Usually this is not in acts of altruism, but by the accumulation of

cash, so in order to get the heroes to hero, you need to place a bounty.

There are two types of bounty. The 'attack' type can be placed on an unfriendly unit or structure. To get the dosh, a hero will have to destroy the

object the bounty was placed on. This is obviously useful for dispatching troublesome raiders or dens of foul creatures too close to the castle. Often, the destruction of a particular building is required to complete a mission.

HOLDING ON FOR A HERO...

There's no shortage of volunteers waiting to take the King's shillings

Heroes come in many shapes and sizes, and unsurprisingly are useful for different things. Rangers have a range attack for example, whereas standard warriors are well armoured and can stand up to a lot of abuse. Gnomes are cheap and dispensable, and useful for exploring. Then there are specialist magic users and the heroes related to the different types of temples you can build. Here is a small sample of some of the characters available.



Warriors

Fine yeomen of stout heart – and deep purses. They are expensive, slow and not too robust until they have risen a few levels.



Rangers

With a ranged attack and a moveable HQ, these units can be quite versatile, but how far can you trust them?



Rogues

If they don't steal the horses, it's because they've eaten them. Petty thieves with a bit of fighting ability – cheap and sneaky



Wizards

The staple of both RPG and RTS games as basic spell casters, the wizards have the best mix and most powerful spells.



Priests

Just like real life, the god-bothersers have various abilities depending on the deity they follow, but they are all magical though.

SYSTEM REQUIREMENTS

Linux Kernel 2.2.x or later
 GNU C Libraries (*glibc*) 2.1.x
 XFree86 3.3.x or later
 166 MHz x86 CPU, or PowerPC G3
 Graphics card capable of 800x600 at 16bpp
 48 MB RAM
 410 MB free hard disk space
 4x CD-ROM drive

In order to destroy something though, you have to be able to find it first. This is where the second type of bounty comes into play. Without prompting, nobody is likely to stray near the dark edges of the map. An exploration flag can be set anywhere though – a hero merely has to visit that spot to collect the loot.

Both of these activities can be an expensive business. The crown's treasury has to cough up at the time the bounty is placed, rather than when the objective is achieved. Bah!

In the counting house

It seems pretty simple. All you need to do as king is explore a bit, set a few

bounties on things and the heroes do the hard work. That is what it comes down to, but that's not all there is to being king. Mostly, the monarch is concerned with his stash of gold. Heroes cost, as do all the buildings they seem to need (Blacksmiths, Inns, Libraries...), and the king is the chump who has to foot the bill. Which in turn means the king's subjects.

The fiscal arrangements aren't too complex though. Tax collectors pop out every now and then to retrieve the king's dues from all buildings. Usually, tax is not enough for a really thriving economy, so specialist buildings such as a marketplace and trading posts can be built. Given that you have to pay to do just about everything, money is important, but build the right things and the pennies will roll in.

For some of the missions, you'll need to build up a substantial battle fund. Magic may be pretty, but it's also expensive. There's nothing more annoying than trying to finish off the last troublesome Troll and discovering that you can't afford to throw any more thunderbolts about.

Because the heroes control their own actions, the king can feel a bit left out of it. For instance, the only way to

try and prioritise exploration is to offer bigger bounties on particular destinations. This obviously gets expensive, and is still no real guarantee as the heroes just bumble along to the most convenient location.

Another drawback to this method is that the heroes associate themselves with one goal, and short of them either achieving it or getting roughed up, they won't consider anything else.

While this does leave the player free to do the deep strategic thinking, it can be a bit tedious at times. However, controlling the heroes individually would be impossible at the pace the game is played anyway.

The Questing beast

So what does the king actually have to do? The different missions cover different types of goals and therefore involve different types of strategy. A few are straight out slug-fests, where you just have to build an army quickly and defend the buildings of the realm. Some require various objects to be sought out, others require particular people to be caught. Several string together different types of goals in the same mission.

Majesty is a fun game, but once the set missions are completed, does it have any longevity? There is the facility to download extra missions, but few of these have actually appeared so far. So the biggest long-term attraction is multiplayer games on the random maps. That said, it will still take a long time to finish the game, and the fact that the maps for each mission are randomised increases the replay value immensely.

There are 17 different missions (some can only be attempted after successfully completing other quests) and it is fair to say that there is a huge amount of variety here. Each is certainly different, and most need to be approached in a different way to win. **LXF**

LINUX FORMAT VERDICT

FEATURES	7/10
PLAYABILITY	9/10
GRAPHICS	8/10
VALUE FOR MONEY	8/10

Not the most thrilling time you've ever had, but a well-crafted game nevertheless.

RATING 8/10



Cash display

The contents of your coffers

Map window

Colour-coded radar blips display the playing area and all visible buildings and characters

Detail panel

Use this to call up detail on buildings, heroes or enemies

Options

Some buildings have other options available, such as recruiting heroes or casting spells

Tracking display

Track a single unit or focus on a building. Click arrows to reset the main view

Scanners

Use buttons to flick through all your buildings, heroes, rewards and enemies



Spell bar

Icons are solid if relevant spell available. Magic-using characters cast their own spells, but you can command a spell (at a price, naturally)

Construction site

Here your toiling peasants are building a new Warriors guild

Guard House

The palace guards rest up in here, and can also fire arrows at marauders from this vantage point

Rogues hideout

Where your Rogue heroes lurk when not stealing anything that's not nailed down

Palace

Your gaff – the place to order construction of buildings, and the repository for your kingly cash

Gnome Hovel

The wee folk aren't pleasant, but they are useful, when they aren't covering indoors when a dragon appears...

House

Your subjects live here. They build the houses themselves too. Hurrah!

Secure Programming Cookbook for C++

Cryptography and authentication has never been so easy for Paul Hudson – or has it?

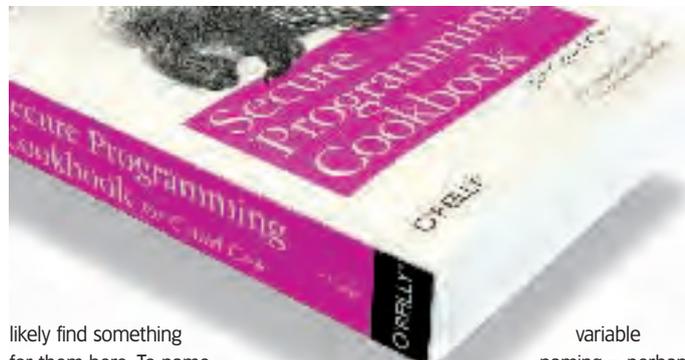


BUYER INFO

- **PUBLISHER** O'Reilly
- **AUTHOR** John Veiga and Matt Messier
- **ISBN** 0-596-00394-3
- **PRICE** £35.50
- **PAGES** 762

Despite what some would have you believe, Linux distros are often just as vulnerable as proprietary OSes. The aim of this chunky tome is to teach more than just how to cope with buffer overruns – instead a more all-encompassing view is taken that starts with explaining each potential vulnerability in your systems, discussing why it's a problem and how it can be abused, and then giving long and well-commented code samples to explain what the code does.

If you're thinking "how can you write 700+ pages about secure programming?", flick through the table of contents – there's so much here on all areas of security that everyone will



likely find something for them here. To name just a few of the highlights, access control is in there, as is input validation, encryption, hashing, PKI, SSL, random numbers, and anti-tampering – each of which are given thorough coverage from authors who clearly have a great deal of experience writing secure code.

Being a cookbook, there are large code samples punctuated with detailed explanations of solutions given. The code is, as usual, designed to be easily dropped into your own apps. Having said that, the authors seem to be quite prone to using somewhat obscure

variable naming – perhaps this is a security measure in itself!

The big downside to this book isn't really its fault, and that's the fact that most people who've been writing C and C++ programs for a few years consider themselves fully knowledgeable and aren't likely to want to admit they need a book on secure programming. This book doesn't make any attempts to order the points, so that real clangers that most of us walk blindly into are mixed in with minor code tweaks. Furthermore, large parts of the writing

build up slowly – you really need to read the chapters in a certain order to get maximum benefit.

If you want to dip your toe in the water, try the free sample chapter at www.oreilly.com/catalog/secureprogckbk/chapter/ch11.pdf. This is a book that's long overdue and makes for an interesting and deeply technical read on a topic that we should all care about more. Yes, it's limited to C and C++ readers, but with the majority of key applications being written in these languages that's where the biggest benefit can be had – give the sample chapter a read, and you'll soon be on your way to the book store to buy the rest of it.

LINUX FORMAT VERDICT

A powerful and initially somewhat scary book that will quickly get you thinking about security while you program – as opposed to as an afterthought.

RATING **10/10**



Practical Programming with TCL/TK

A thick guide to a fascinating system, but can it TCL Richard Cobbett's interest?

BUYER INFO

- **PUBLISHER** Prentice Hall
- **AUTHOR** Brent B. Welch, Ken Jones, Jeffrey Hobbs
- **ISBN** 0-13-038560-3
- **PRICE** \$49.99
- **PAGES** 650

TCL is an interesting language. While intended to be easy to use and pick up, it is packed to bursting point with little eccentricities – ranging from the use of C/C++ and most recently Java to extend its extremely small core command set, to the way that even commenting out a command won't prevent its contents from being processed the next time that you run the program. This book is an exceptionally easy way to start out in this language, although it must be said that some prior coding knowledge is an advantage when working through it. Both of the main sections – the first



third of the book dealing with the Tool Command Language itself, moving onto creating windowed applications in the second third – begin simply enough with Hello World, but it's not long before other languages are brought into play, with the first really meaty example coming in the form of a simple CGI guestbook.

Nevertheless, the simplicity of TCL itself ensures that you should be up-

to-speed with at least its part of the process by that stage. It's always good to see a language that can be summarised in a couple of sheets, and you'll soon get used to seeing the standard subjects in each subsequent example. The section on TK also does a sterling job of actually showing you what your code snippets are geared towards creating – with coherent mini-examples of everything

from internal padding to simple horizontal fills.

While the book, like TCL, is platform-agnostic, it goes into plenty of depth on getting your software running on each platform – right down to explaining the differences between different versions of the packages (both in the sense of 'Wish.exe' vs 'wish' vs 'Wish', and a more technical look at exactly what version 8.3 added to TCL itself). While of course it can't make you a good enough C/Java programmer to extend TCL every way you might want, it can at least help you prepare your scripts so that they're ready to face such extension.

LINUX FORMAT VERDICT

An excellent way to get scripting, with the power to potentially produce even more impressive results.

RATING **8/10**



Intrusion Detection with Snort

Intrusion detection is a black art to most of us – Paul Hudson wonders whether this book change that...

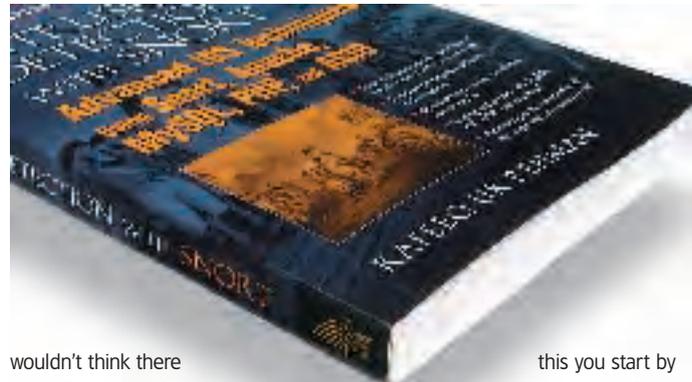
BUYER INFO

- **PUBLISHER** Prentice Hall
- **AUTHOR** Rafeeq Ur Rehman
- **ISBN** 0-13-140733-3
- **PRICE** \$39.99
- **PAGES** 263

Intrusion detection has been covered many times before, but often only from an abstract, "here's what it is" perspective. The few books that cover *Snort* specifically are often so focused on the topic that you need to buy a couple of *other* books to figure out what you can do with your *Snort* information. For an interview with *Snort*'s creator, Marty Roesch, see *LXF43's Linux Pro*.

This book, however, has the subtitle *Advanced IDS Techniques using Snort, Apache, MySQL, PHP, and ACID*, so it's quite clear that this volume will be a solution-driven piece even before you open its cover.

At only 263 pages including the index and table of contents, you



wouldn't think there was much scope for covering so many topics in depth, but a brief flick through reveals that the non-*Snort* topics are covered only as far as they relate to *Snort*. For example, installing MySQL is covered in just half a page, whereas actually creating and using MySQL tables to store *Snort* data is twenty pages.

The goal of the book is to take you step by step through building up a complete IDS system that is accessible over the web. To facilitate

this you start by just setting up *Snort* (to which, unsurprisingly, the largest chunk of the book is devoted) and reading your info from logs. You then move onto storing your IDS data in a MySQL database, allowing you a little more flexibility than just straight logs. Finally the book walks you through installing and configuring ACID, the Analysis Console for Intrusion Databases, which allows you to analyse the *Snort* results that get stored in the MySQL tables.

The flow of the book is very much like a series of step-by-step tutorials, which means it's easy to stop and start as necessary and/or skip ahead if you know a particular section already. On the down-side, there are far too few illustrations used to break up what is otherwise a very heavy topic that is likely to bore people who don't have a particularly strong interest in the area. But, what text there is is well written and to the point, if a little dry. While it's certainly not the most comprehensive book on *Snort* or MySQL, it's one of the few that manages to bring the two technologies together and do so in under a thousand pages.

LINUX FORMAT VERDICT

A bit of a slow read and it's certainly not an exciting topic, but the author does his best to make a rather obtuse subject somewhat more interesting.

RATING **8/10**



Who will keep you on the Information Highway?

To date **Hostway** services the hosting needs of over 300,000 customers worldwide, facilitating millions of dollars in e-commerce transactions every month. Fully operational centres in Chicago, London, Seoul, Amsterdam and Sydney, underline our unique position as the only truly global hosting provider.

With a core competency in Linux based hosting since 1996, **Hostway** uses the open-source platform for many of its own turnkey systems. Over the years, this has helped us focus on optimising the benefits that our Linux users enjoy - ease of use, maximum stability, and high-specification.

We are consistently at the forefront of research & design in Linux technologies, and are firmly committed to leveraging the LAMP (Linux/Apache/MySQL/PHP) platform with a RedHat back-end to deliver highly stable, flexible, yet accessible hosting solutions.

Hostway - straightforward Linux hosting.



A FREE MONTH
of web hosting.
When ordering online
use "FUTLINO48".

	KEY FEATURES	DISK SPACE	MONTHLY TRAFFIC (Gb)	SHELL-ACCESS	PHP4.3.2 SCRIPTING	MYSQL3.23 DATA AREA	ADDITIONAL FEATURES	ACTIVE CGI LIBRARY	CRONJOBS	XML SUPPORT	PYTHON	FRONTPAGE EXTENSIONS	POP3/SMTP	EMAIL ACCOUNTS	UNLIMITED ALIASES	EMAIL FORWARDING ALIASES	MEDIA FEATURES	HTTP STREAMING	REAL SERVER PROFESSIONAL	MONTHLY FEE*	YEARLY FEE*	SETUP FEE*
LINUX SILVER	100	1	Yes	Yes	0	Yes	Yes	Yes	Yes	Yes	Yes	5	Yes	Yes	Yes	No	Yes	No	£8.95	£89.95	FREE	
LINUX GOLD	200	4	Yes	Yes	1	Yes	Yes	Yes	Yes	Yes	Yes	10	Yes	Yes	Yes	No	Yes	No	£14.95	£134.95	FREE	
LINUX GOLD PLUS	200	6	Yes	Yes	1	Yes	Yes	Yes	Yes	Yes	Yes	30	Yes	Yes	Yes	No	Yes	No	£19.95	£179.95	FREE	
LINUX PLATINUM	500	10	Yes	Yes	2	Yes	Yes	Yes	Yes	Yes	Yes	50	Yes	Yes	Yes	No	Yes	No	£49.95	£449.95	FREE	
LINUX PLATINUM PLUS	500	15	Yes	Yes	2	Yes	Yes	Yes	Yes	Yes	Yes	100	Yes	Yes	Yes	Yes	Yes	No	£89.95	£809.95	FREE	

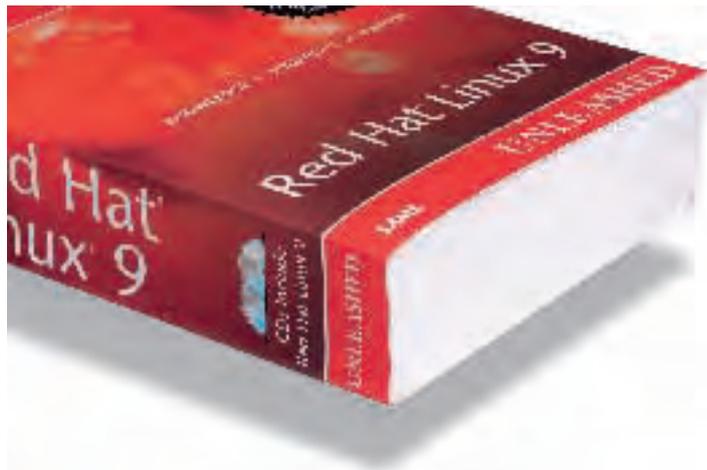
All plans come with a 30 day money-back guarantee | * excluding VAT & subject to change.

web hosting • domain registrations • dedicated servers • resellers 0808 180 1880 • sales@hostway-uk.com • www.hostway-uk.com

HOSTWAY
THE HOSTING COMPANY

Red Hat Linux 9 Unleashed versus Red Hat Linux 9 Professional Secrets

Amias Channer referees the battle-royale of the heavyweight Red Hat Linux 9 guides.



BUYER INFO

RED HAT LINUX 9 UNLEASHED

- PUBLISHER SAMS
- AUTHOR Bill Ball and Hoyt Duff
- ISBN 0-672-32588-8
- PRICE £36.50
- PAGES 1056

RED HAT LINUX 9 PROFESSIONAL SECRETS

- PUBLISHER Wiley
- AUTHOR Naba Barkakati
- ISBN 0-7645-4133-1
- PRICE £34.95
- PAGES 1038

Red Hat's recent announcement that version 9 will be its last non-enterprise release and that the future of its non-enterprise distro is as part of the Fedora project creates an interesting situation for the Red Hat 9 guidebook market.

Whilst good sysadmin skills are definitely helpful for home and small-scale networks, the money to pay for the book is probably less forthcoming. For professional users, Red Hat 9 will probably not be on the menu (in favour of an Enterprise offering). With the Enterprise platform using a modified Red Hat 7.x base so many of the advanced sysadmin topics may be redundant. Fedora will start with the Red Hat 9 core but is unlikely to retain compatibility for long. However, with the first Fedora release still around 6

months away, obsolescence for Red Hat 9 is still distant enough to keep it viable and make one of these books a worthwhile purchase.

Red Hat Unleashed, co-authored by long-time LXF contributor Hoyt Duff, is a catch-all kind of guide. It reads in a manner that would be suitable for someone completely unused to Linux, but is simultaneously not too patronising for experienced readers, commendably fixing what some considered to be the main shortcoming of *Red Hat 8 Unleashed*. With a title like 'unleashed', one might expect a guide to 'getting the most' from an existing system, but the focus is much less specific. With over a thousand pages in 27 chapters spread over 5 parts, a very comprehensive index and effective round-ups of the topic at the end of each chapter, this is quite an epic volume. All common aspects of Linux use are well covered, from services like HTTP and Samba, to multimedia applications and hardware configuration. Unfortunately, a number of errors appear to have crept into the grammar, and confusingly on a couple of occasions, the commands in the walkthroughs were invalid due to missing spaces. This would not be a problem for readers with more than a passing familiarity with Linux, but since the emphasis of the book is not on the command line, novices could be thrown by this typographical problem.

The experience and knowledge of the authors shows heavily in the

abundant tips and notes boxes where most of the common Red Hat pitfalls are worked around, good sysadmin practice is expounded and worthwhile advice on all aspects of deployment and suitability is offered.

Professional secrets

Red Hat Linux 9 Professional Secrets is targeted more toward the hacker mindset than *Unleashed's* very commercial feel. As a result, this book rarely asks you to consider the implications, assuming that you know why you are doing this already – a slightly dangerous approach, but a little more exciting to read. Given Red Hat 9's very complete GUI, the author's choice to use the command line as much as possible is going to split the audience. However, a useful side-effect of this approach is that quite a lot of the examples will apply to other Linuxes which may prove very useful to users of multiple distros; as will the frequent pointers to relevant RFC documents. The layout of this book is well thought-out with clear tables and informative 'secrets' boxes that make the content easily accessible; but having a much smaller index makes the book harder to dip in to.

In most of the chapters there is plenty of documentation for configuration file formats, which often gives you access to more complex and powerful configuration options not present in some GUIs.

Although the lack of documentation for either of the databases shipped with Red Hat 9 seems a little short-sighted, this is an excellent and concise guide.

Unleashed (despite typos) feels like a more polished product, and its style will suit the more business-oriented user, whereas *Professional Secrets* is slightly less slick in the delivery and print quality. The inclusion of large sections on programming in both books seems a little unnecessary – to be properly usable, much more information is needed and better language-specific guides already exist. *Professional Secrets* is poorer in this respect, with a huge java programming section in the middle of 200 pages on programming.

Both books come with the publishers' edition of Red Hat 9 on 2 CDs to help get you started. [LXF](#)

LINUX FORMAT VERDICT

Either of these books will get you up-to-speed with Red Hat 9 very easily, the decision really comes down to which interface you prefer, *Professional Secrets* for the command line or *Unleashed* for using a GUI.

RED HAT 9 UNLEASHED
RATING **8/10**



RED HAT LINUX 9 PROFESSIONAL SECRETS
RATING **9/10**



Roundup

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



OUR SELECTION AT A GLANCE

- Brahms
- MuSE
- RoseGarden4
- Jazz++
- Anthem

MIDI sequencers

Richard Drummond puts five MIDI packages to the test to find out how far Linux has to go in order to cut the mustard as a professional audio platform.

The history of music production on desktop computers dates right back to 1975 when Steve Dompier found out that the radio interference generated by the first desktop machine – the Altair – could be controlled sufficiently to induce something approaching music from a nearby loud speaker. Thankfully, technology has moved on somewhat since then! In fact, your average PC has more than enough audio power to be used as a complete home recording studio – you just need the right software and a bit of talent to make it happen.

On Windows and Mac OS, popular music production suites include *Cubase* and *Cakewalk*. But what about Linux? Are there any Open Source packages to rival the functionality of these commercial heavyweights? As it happens, there are – and we're going to examine five of them right here.

The applications on test are first and foremost MIDI sequencers; MIDI (Musical Instrument Digital Interface) being the universal language spoken by electronic musical instruments. In essence, sequencers allow you to compose, arrange and manipulate

sequences of MIDI events and output them to a MIDI device. That device may be an external MIDI instrument such as a keyboard, a hardware synthesiser built into your PC's sound card, or a virtual MIDI synthesiser created in software. It's all the same to the sequencer. All that differs really is the sound quality. Increasingly however, these types of applications also allow audio sequencing – that is, the playing of digital audio samples can be synchronised with your MIDI tracks – and our Linux packages are no different.

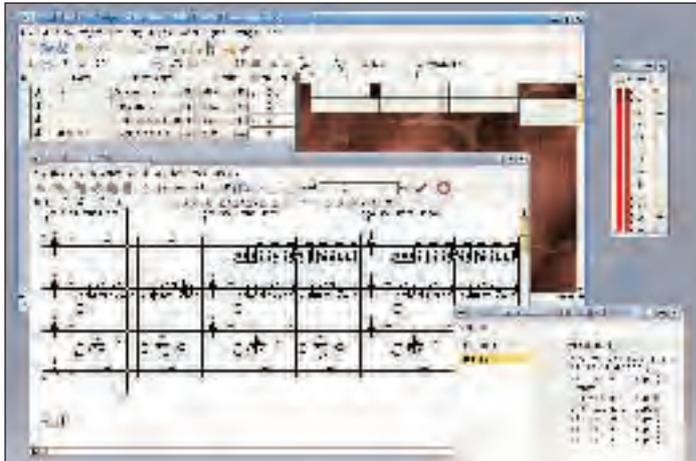
The MIDI sequencers we are looking at vary greatly in the range of features they offer, but all employ a broadly similar user interface model. Disregarding a few hiccups, they are all fairly easy to use. What may be less easy is installing and configuring the underlying audio platform that these programs require to run.

We're at a difficult time for Linux audio at the moment. The new Linux sound API – ALSA – is becoming more widely adopted, but it can be problematic to configure and driver support may be still be shaky depending on your hardware. What's more, there's a seemingly endless stream of audio-related libraries and toolkits to further bamboozle the user: *aRts*, *JACK*, *LADSPA* (see the *Acronym overload* box on page 36). When these technologies find their way into your average distro in stable form, expect life to become much easier.

Brahms

Simple, aRts-based MIDI application.

■ **VERSION** 1.0.2 ■ **WEB** brahms.sourceforge.net



Thanks to the aRts MIDI manager, *Brahms* is easy-to-use, but it lacks professional features.

The aRts sound system, as popularised by the KDE desktop, has its detractors, but is widely regarded as a powerful audio framework. What is less well known, perhaps, is its MIDI capability.

With the aRts MIDI Manager, users can simply hook-up physical and virtual MIDI devices to applications and make use of the aRts software synthesiser as a MIDI output – all at the click of a button.

Brahms supports MIDI and audio sequencing, and can output to ALSA as well as aRts. The choice between which system to use – aRts or ALSA – is automatically made at start-up depending on whether the aRts daemon is running, and you cannot change systems on the fly. *Brahms* provides few controls for configuring output itself. In aRts mode, you may only globally assign a MIDI port to *Brahms* in the MIDI Manager; you may not target different devices for different tracks like you can in ALSA mode.

The *Brahms* user interface is adequate, but can be counter-intuitive in areas and lacks polish. Tool-tips help navigation somewhat and the docs are useful, though incomplete. *Brahms* employs the usual track/part model for structuring songs, albeit with few editing controls, and offers the standard score, piano-roll and event editors for editing parts. Recording from a MIDI device is not supported, so you must input tracks with the mouse and keyboard. The score editor makes this moderately painless due to some handy keyboard shortcuts.

Transparent software synthesis thanks to aRts means it's easy to use even without MIDI hardware. It's not really suitable for professional work,

though. The lack of MIDI input and the clunkiness of the GUI mean it's less pleasant to use than it should be. These problems could be overlooked, but other more serious issues cannot be.

Firstly, *Brahms* hasn't seen any development in the last year. The current release will build against KDE 2.x or 3.x, but only supports the obsolete ALSA 0.5. A development version targeting ALSA 0.9 may be downloaded, but this requires some hacking just to compile. Secondly, *Brahms* lacks stability. Some of the blame may be laid at the feet of aRts, and some on ALSA, but *Brahms* just isn't as robust as it should be. The MIDI import function can be particularly problematic.

LINUX FORMAT VERDICT

FEATURES	7/10
EASE OF USE	7/10
DOCUMENTATION	6/10
PERFORMANCE	5/10

Potentially a simple but capable sequencer, but marred by a lack of stability and GUI polish.

RATING **6/10**



MusE

The band had a hit with 'Falling down' – will this?

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



MusE has a pleasant GUI and some powerful features, but still needs work.

MusE is a Qt-based music editor that is fully buzz-phrase compliant. It supports sequencing of MIDI and audio tracks, and can do MIDI output to OSS,

ALSA and serial ports and audio output to ALSA or the JACK system. It features a number of plug-in software synthesisers (including support for

fluidsynth). It implements a LADSPA host, enabling the use of a wide range of plug-in audio processors and effects and uses the LADCCA system to load and save audio application sessions.

MusE is actually quite pleasant to use. Setting up all these various pieces of technology is somewhat less so – but thankfully optional. *MusE*'s GUI is clean and nimble, but not without quirks: for instance, floating toolbars seem to lead a life of their own sometimes and don't stay where they're put. The docs are meagre, but cover enough of the basics to get you started. The familiar track and part concept for arranging a song is straightforward and flexible – you can create parts, cut, paste and drag them. Neat touches to help navigation include arbitrary naming and colourisation.

Composing parts is no problem either, and *MusE* is equipped with the usual range of editors. The score editor is particularly good, producing great-looking scores onscreen or as hard copy, but its range of supported musical marks isn't as broad as some others. I do like the way that the score editor renders the staff over multiple rows rather than as one single long staff

across the screen. Recording from a MIDI input is supported, but keyboard control in the score editor is sadly non-existent. Support for audio tracks is still new in *MusE* and the editor for sequencing audio samples is spartan.

MusE boasts some professional features, including low latency audio output via JACK and support for MIDI synchronisation (it can act as a master or slave). However, like many of the sequencers here, it is beset by bugs. During testing I was unable to get any of the software synths to work at all, and sequencing audio tracks caused segfaults more often than not.

LINUX FORMAT VERDICT

FEATURES	8/10
EASE OF USE	7/10
DOCUMENTATION	6/10
PERFORMANCE	7/10

A slick sequencer with some professional features but also plenty of bugs, alas. Nevertheless, *MusE* shows a lot of promise and at least it's still being developed.

RATING **7/10**



Rosegarden4

The closest thing to Cubase for Linux.

■ **VERSION** 0.9.1 ■ **WEB** www.all-day-breakfast.com/rosegarden/

Steinberg, creator of the popular Cubase range of music production software on Windows and Mac, unsurprisingly don't support Linux; however, it has been said that *Rosegarden4* is the closest thing to a *Cubase* for Linux yet available – and it's easy to see why.

Rosegarden4 is an incredibly powerful MIDI and audio sequencer that is so laden with features it puts some of the other apps on test here to shame. It is built on the KDE3 platform and supports either ALSA or *aRts* for MIDI and audio I/O (you must choose one or the other: it cannot be built to support both), JACK for audio and LADSPA for audio plug-ins. I had problems compiling it with GCC 3.2 against KDE 3.1.4, but pre-compiled binaries for many of the leading distros are available for download.

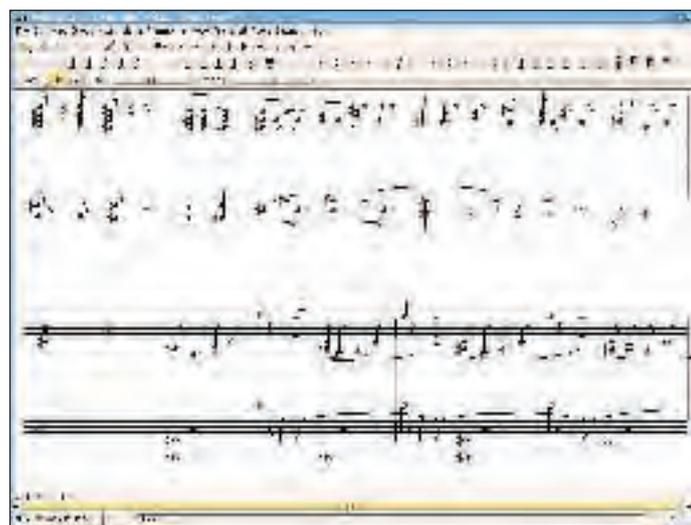
For older systems, *Rosegarden4's* ancestor, *Rosegarden2.1* is still available for download. This uses plain X for the GUI, is MIDI only, and only supports OSS for I/O, but it's considerably lighter than the new version and more stable. *Rosegarden4* itself is still in the beta-testing phase, with a version 1.0 release expected soon. While it's less buggy than some of its rivals, I found *Rosegarden4* still apt to die inexplicably when you try to do something that it doesn't like.

Along with the sunshine

Rosegarden4 employs the same track and part model for arranging music as the other applications on test – although it calls parts 'segments'. I found its tools for manipulating tracks rather cumbersome and unintuitive – particularly when trying to map instruments to tracks – but its part editor is second-to-none. Its easy to 'draw' segments on screen, and to drag, resize, rescale or cut and paste them. A particular highlight for me was the tools which split up parts on silence or pitch, which are really handy when you want to edit music from an imported MIDI file.

The usual score editor (which *Rosegarden4* calls the 'notation editor'), piano-roll editor (the 'matrix editor'), event list editors are included – but no drum editor. Recording from a MIDI device is supported, but if you don't have a piano keyboard, both the notation and matrix editors are supremely usable just with mouse and keyboard thanks to a plentiful supply of keyboard shortcuts.

The notation editor in particular is excellent, and produces beautiful-looking scores on screen. It also supports output via the music typesetting package *Lilypond*, so even



The notation editor is easy and quick, and can produce beautiful scores.

more beautiful hard copy is possible. The notation editor is a delight to use and just packed with features. It supports linear or page-based layouts, and a great touch is that you can display segments from multiple tracks in a single notation editor – they'll be displayed on multiple staves. A comprehensive set of musical notation and marks is supported including slurs, tuplets, grace notes, accents and more. Tools are included for manipulating a staff's clef, and to perform key and time signature changes. In fact, the notation editor makes *Rosegarden4* a capable score editor – even if you're not interested in all that MIDI nonsense.

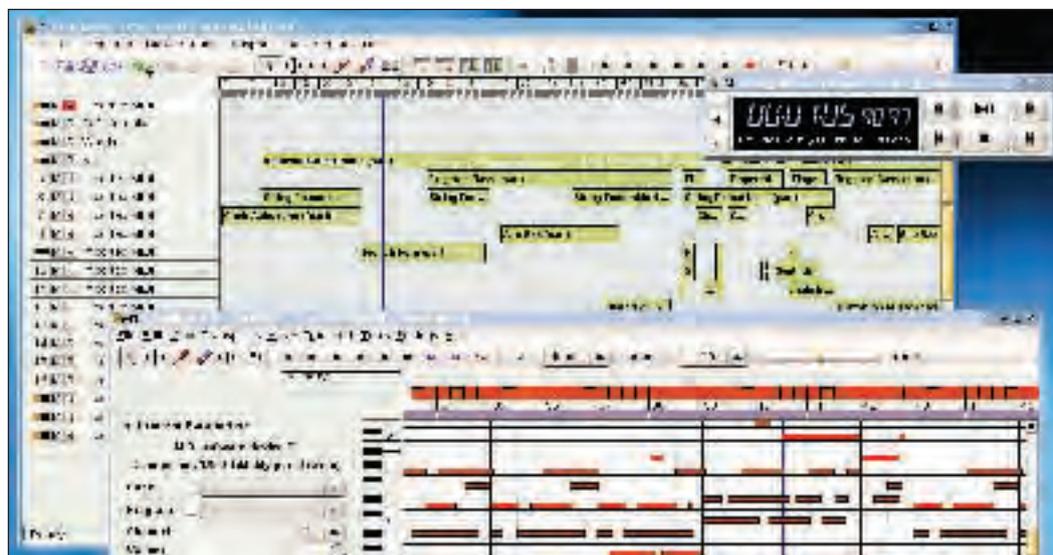
A little rain sometimes

Rosegarden4 is good, but not perfect by any means. There are still plenty of

bugs, and I could not get audio sequencing to work at all.

Rosegarden4 is very configurable, but the price is that it's not always easy to configure. The documentation supplied is informative, but has some great gaping holes. Thankfully, there are a couple of tutorials available on the web, which should help newcomers. Generally, *Rosegarden4* is easy to use, and tooltips and help bubbles aid immensely. However, I am not fond of several areas of the GUI layout, including the track listing and the segment parameters and instrument parameters docks. I would much prefer to be able to tear these off as separate windows. And, while the dial widgets might look cool, a standard spinner would be much easier to use. Finally, *Rosegarden4* can be more than a little sluggish in operation, especially the notation editor.

Despite these irks, *Rosegarden4* is without doubt the most complete MIDI package for Linux. What's more, it's still actively developed, unlike some of the other applications on test. That fact alone should make it the most sensible choice for many users.



Rosegarden4 creates a powerful environment for sequencing and editing music.

LINUX FORMAT VERDICT

FEATURES	9/10
EASE OF USE	7/10
DOCUMENTATION	6/10
PERFORMANCE	8/10

Clearly the Rolls Royce of music production packages on Linux, but not without rust spots here and there.

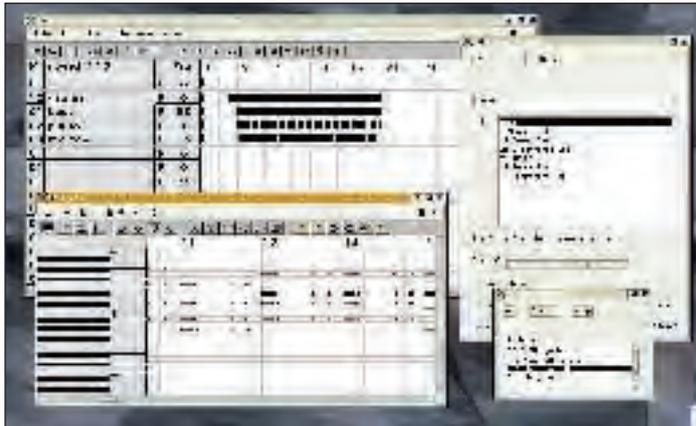
RATING **8/10**



Jazz++

You don't have to be 'Bleeding Gums' Murphy...

■ **VERSION** 4.1.8 ■ **WEB** www.jazzware.com/cgi-bin/Zope.cgi/jazzware/



You can patch *Jazz++* for ALSA 0.9 support so it'll work on modern systems, but I wonder if it's really worth the effort?

Hands up who knew there was once a commercial sequencer for Linux? Anybody? I didn't either. Well there was: Jazzware's *Jazz++*. It's now an Open Source project, but it hasn't

been developed much in the last three years – which is a shame because, despite an unattractive GUI, it's a fairly powerful application with some nifty advanced features.

Jazz++ supports sequencing of MIDI and audio tracks, and can output either to OSS devices or just MIDI to ALSA. However, it is designed to work with the now obsolete ALSA 0.5, so that won't work on modern systems. A patch is available to implement ALSA 0.9 support, but the lengths you have to go to compile *Jazz++* means that few will bother. You see, *Jazz++* uses an old version of the wxWindows toolkit for its GUI (version 1.68) which – even if you can find a copy – won't compile easily on recent distros, and *Jazz++* won't work with newer releases of wxWindows. *Jazz++* also includes its own kernel driver for real MPU401 devices, but it's provided as a patch against a 2.0 kernel, so isn't particularly useful.

The old wxWin toolkit used here is built on *Motif* and hence looks and feels dated. If you can get passed that, then *Jazz++*'s interface isn't too bad. It lacks a score editor, but otherwise it is bristling with features, including neat things like a rhythm generator (for creating natural-sounding drum grooves), an arpeggio generator and a harmony generator. Control over MIDI output, instrument mapping, etc. is very flexible. Ease of use is hindered by

Jazz++'s small, unintelligible toolbar icons and a lack of tooltips. Online documentation is provided, though.

If you can get *Jazz++* to compile and run on your system, then it is powerful and quite usable. In fact, during testing it was by far the most robust of the applications being reviewed here. Moreover, there's a wealth of documentation, resources, and examples available at the Jazz web site, so *Jazz++* should prove easy to learn. I just wish somebody would pick up development of the project, starting first with a port to wxWin 2.x. Otherwise there's a lot of hard work and potential going to waste here.

LINUX FORMAT VERDICT

FEATURES	8/10
EASE OF USE	4/10
DOCUMENTATION	7/10
PERFORMANCE	6/10

With a dated – but surprisingly usable – interface, this app's most significant problem is that it so difficult to compile.

RATING **6/10**



Anthem

A second sequencer based on a commercial app.

■ **VERSION** 0.0.17 ■ **WEB** <http://anthem.sourceforge.net/>



Anthem doesn't want to work well on KDE 3.1. Note the empty window where the piano-roll editor should be, and the missing icons.

Like *Jazz++*, *Anthem* is derived from a commercial product; but for RiscOS rather than Linux. The Open Source *Anthem* looks very similar to Trax Software's commercial version, but it is now built on the KDE framework.

The sequencer engine of *Anthem* has been spun off as a separate project – *TSE3*, see <http://tse3.sf.net>. It's mature, portable and has a simple C++ API. On Linux, it supports output to MIDI devices via OSS, ALSA and aRts APIs.

Like many of the apps on test here, *Anthem* is suffering from a lack of development: the last release was 16 months ago. It should build on either KDE 2.x or 3.x, but it doesn't seem to support newer KDE 3.x releases at all well – and it needs modification to compile with GCC 3. Should you manage to build *Anthem* and get it working, then potentially you'll have a fairly sophisticated MIDI sequencer. Thanks to *TSE3*, configuring MIDI output is a breeze and there are plenty of controls, such as MIDI filtering, echoing, and so on. *Anthem* uses a more object-based model than the other apps. For instance, you can set MIDI properties globally for a song, and locally for each track or track part. By default, a track inherits the song's properties and a part inherits its track's properties.

Potentially, *Anthem* should provide a comfortable and intuitive editing environment – despite an utter lack of documentation – but there were too many features that just did not work for me, perhaps due to incompatibilities with KDE 3.1 (you may do better with KDE 2.2, but we didn't test it with that). Colourisation of parts in the part view did not work, which – thanks to the

(non-modifiable) backfill image used – rendered part names illegible. Worse, the phrase editor and piano roll editor would not even appear at all. Combine that with an all round lack of stability, and your left with an exciting but just plain non-functional app. This a shame because I know there is lots of good stuff in here. *Anthem* is very configurable, has a great GUI (the bits that work!), and has tons of useful features.

Anthem is supposed to a work-in-progress, but work isn't progressing. If it were being actively developed you could forgive such problems, but unless it's resurrected, it remains an interesting but almost useless piece of technology.

LINUX FORMAT VERDICT

FEATURES	7/10
EASE OF USE	6/10
DOCUMENTATION	2/10
PERFORMANCE	4/10

Has all the makings of a great MIDI sequencer, but a lack of development has left too many problems for it to be anything more than a curiosity.

RATING **4/10**



MIDI SEQUENCERS THE VERDICT

After having a look at these packages, what's the conclusion? How does Linux stack up as a MIDI sequencing platform? The answer has to be not very well. Let me explain.

The crucial issue is ease of use. *Not one* of the eight or nine MIDI sequencers I tested while writing this article would compile on my system out of the box without some modification and some, particularly *Jazz++*, required a lot of grappling with libraries and toolkits. This is not what the end user should expect, certainly not a user who is just interested in writing music on their Linux box. But since few of these packages provide pre-compiled binaries it is a problem that has to be faced.

To be fair, many of the ease-of-use problems with these packages can be blamed on underlying audio platform and the difficulties in configuring *aRts*, ALSA or whatever. All of these sequencers are development versions, not stable releases, so problems are to be expected, and even ALSA is still arguably in a beta state. Of course,



fewer users are interested in MIDI, so these aspects are less well tested and debugged. However, several of these applications suffer from a critical lack of development, and so poor support for modern systems is just one of the consequences.

The second significant problem with all these packages is stability. With the exception of *Jazz++*, none are robust enough for production work.

There's nothing like a crash that wipes out hours of work to frustrate the end user (although, at least *RoseGarden4* has an autosave feature).

Best of the bunch

They all have problems, but which is them of them is most functional? Regular readers may notice that this page lacks the usual table comparing features available. We spent a lot of time trying to construct one, but in the end had to admit defeat, as there were significant differences in most of the apps reviewed here between functions that were *claimed* to be supported, and what *actually worked* for us.

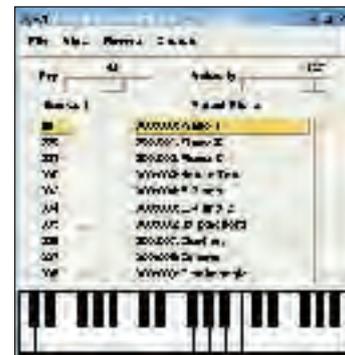
Rosegarden4 stands head and shoulders above the others in terms of features. As has been said, it is the nearest to *Cubase* for Linux that is yet available. When it stabilises, *Rosegarden4* will be quite suitable for professional use; but there is still a lot of work to do yet, and it could still benefit from more polish and certainly more documentation.

How about the others? Many of the others show a lot of potential, but more work is needed. I liked *Anthem*, but it's just not usable as is. *MusE* is fast and has a good range of features, but is way too buggy for real-world use. *Brahms* similarly has too many stability issues. In fact, despite its ugly GUI, I would have to rate *Jazz++* in second place to *Rosegarden4*.

Nevertheless, I am disappointed with the current state of MIDI sequencing software on Linux. The only solution is if more users



RoseGarden4 has an impressive array of features and is the most professional of the sequencers here.



If you lack of real MIDI keyboard, then a software keyboard such as *vkeybd* makes entering music faster.

contribute to these projects. If you are looking for an Open Source project to adopt, consider one of the MIDI sequencers. They all could benefit from a lot more development, testing and the reporting of bugs. [LXF](#)

ACRONYM OVERLOAD?

Untangle some of the Linux audio jargon

MIDI (Musical Instrument Digital Interface)

The ubiquitous hardware and software standard for getting electronic musical instruments to talk to each other. See www.midi.com/ and www.midi.org/OSS (Open Sound System)

The existent sound driver framework for Linux and many other Unices, which supports audio and MIDI devices. www.opensound.com/

ALSA (Advanced Linux Sound Architecture)

A replacement for OSS designed to address many of its shortcomings. It will be the default sound API with Kernel 2.6 and is already shipped with many leading distros. See www.alsa-project.org/

aRts (Analog Realtime Synthesiser)

Originally a toolkit to emulate an modular analog synthesiser, this has evolved into a powerful media streaming framework and is employed as such by the KDE desktop. www.arts-project.org/

JACK (Jack Audio Connection Kit)

A low latency audio server that allows applications to share an audio device or to stream audio to one another. See <http://jackit.sourceforge.net/>

LADSPA (Linux Audio Developer's Simple Plugin API)

A system that gives audio applications access to a library of re-usable plug-in components to perform audio processing and effects. See the website for more details: www.ladspa.org/

LADDDCA (Linux Audio Developer's Configuration and Connection API)

A system for performing session management with ALSA and JACK audio applications. See <http://pk1.net/~node/ladcca.html>

AND THERE'S MORE

Some extra sequencers to try

We didn't have space to cover all Linux MIDI sequencers that are available, but there are a couple more. They all suffer from comparable stability and under-development problems to the ones on test, except for *Tukta*, as featured on page 40 of this issue.

SHAKE TRACKER

www.reduz.com.ar/cheesetronic/

Dares to be different by providing a sound tracker style interface to MIDI sequencing. I had lots of problems building it, though.

MIDI MOUNTAIN

www.midimountain.com/

Looks promising and uses GTK+ rather than Qt or KDE for its GUI, but hasn't been developed for over two years.

BINARS

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

Based on GNOME, but again hasn't been updated for over two years.

Hot Picks

The best new open source software on the planet!



Mike Saunders

A coder since Amiga times, Mike's a Linux and BSD guru.

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

HOT PICKS AT A GLANCE

DOSBox	38
Firestarter	39
GAIM	39
CUTE	40
Tutka	40
AmaroK	41
Netwox	41
LTris	42
Atomic Tanks	42

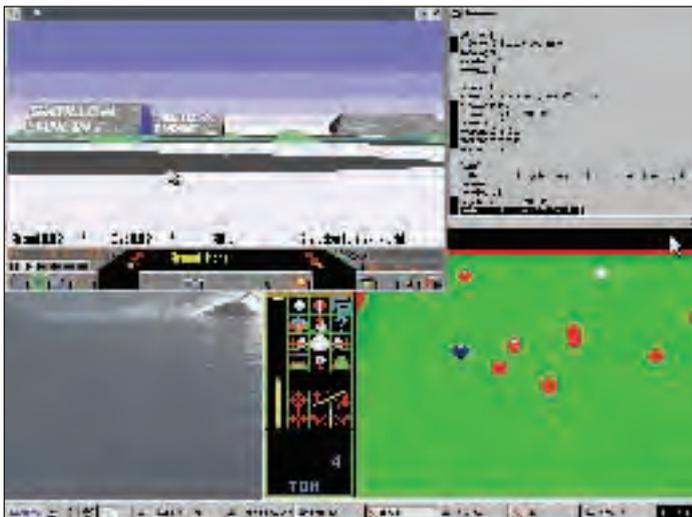
HOT PICKS AWARD

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!



PC & DOS EMULATOR DOSBox

■ VERSION 0.58 ■ WEB <http://dosbox.sourceforge.net>



Frontier, Microprose's *F1GP* and the config file. Running two games at once is possible, but rather sloooooow...

Progression in the computing world

– for all the benefits of new technologies and performance improvements it brings – suffers from a big problem: retaining backwards compatibility. There's always a trade-off between incorporating the latest goodies while still allowing old software to run, and it's often a headache. This isn't such a big problem for Linux; most Open Source software can be tweaked and recompiled for newer distros, but on the Windows side, Microsoft's switch to the NT kernel for XP means that many old DOS programs won't run.

To combat this problem, *DOSBox* was born. It's essentially a PC emulator with a built-in version of DOS. Luckily for us, it's not just limited to the Win32 platform and users of alternative systems can re-live the old days of VGA games and tweaking CONFIG.SYS. Seeing as most popular application software has had newer versions released since DOS days, the coders behind *DOSBox* have concentrated on game compatibility instead.

If you've been on the Linux scene for some time, you might be thinking 'What happened to DOSEMU?' A valid question – at first glance, both tools appear to have the same goal. The most crucial difference, though, is that *DOSBox* provides a full CPU emulator whereas DOSEMU uses x86 'virtualisation' tricks (as with *VMware*), and the result is that DOSEMU only works on standard PCs whereas *DOSBox* can run on Apple Macs, Sparcs and Alphas etc.

As this issue was being prepared, *DOSBox* was at version 0.58 with a new release in the pipeline. We've included the most recent official source on the coverdisc and a working CVS snapshot – the latter is particularly useful for running protected mode games. If that's not important, it's best to stick with the stable releases for now. Thankfully, *DOSBox*'s only major dependency is SDL (for the front end), and that should be installed on most distros or can be found on our CD.

Box clever

When first run, *DOSBox* throws up a familiar text screen (similar to the Windows DOS/CMD prompt) with a handful of useful commands thrown on the Z: drive. Initially, the most important of these are CONFIG.COM, which can be used to write out a *dosbox.conf* file containing all the settings, and MOUNT.COM for attaching Linux directories as drives. This works in a similar fashion to DOSEMU's LREDIR util, and means no messy partitioning is necessary.

With an appropriate directory mounted, you simply need to 'cd' into it and run the game. Satisfyingly easy and hassle-free! Whether or not the game will work is another matter, but the compatibility was applaudably good in our tests – *F1GP*, *Frontier* and *Jimmy White's Snooker* all worked decently. *Worms*, a protected mode game, ran very slowly; hopefully the newer release will sort out the speed issues. A useful addition is the ALT+Enter combo for a fullscreen mode.

DOSBox's config file is straightforward and each option is self-explanatory. Frame skipping and cycles per loop can be modified to ensure an accurate game speed, while SoundBlaster emulation parameters (port, IRQ etc.) can also be altered. Overall performance doesn't match DOSEMU, which is inevitable with a fully emulated CPU core, but it's vastly more simple to set up and pretty much works straight out the box. DOSEMU's development limbo makes *DOSBox* a better choice for fixes and support too, and while *Bochs* does a commendable job of supporting multiple OSes, it's naturally not as strong as *DOSBox* when it comes to a pure gaming focus.

With proper protected mode support due in future releases, *DOSBox* is establishing itself as a must-have tool for 1990s PC gaming fans. We're hugely impressed by its ability to work without hours of tweaking, and coupled with the wide compatibility range and acceptable performance even on a 500MHz box, it's one to watch and another feather in Linux's desktop hat.

FIREWALL TOOL

Firestarter

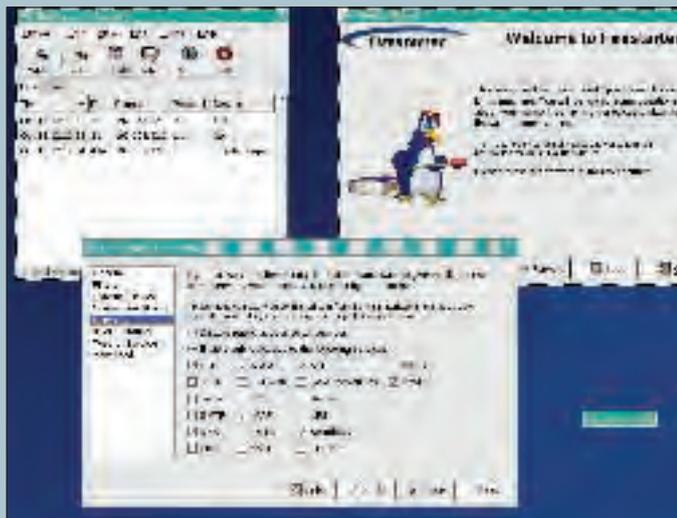
■ VERSION 0.9.2 ■ WEB <http://firestarter.sourceforge.net>

Keith Flint may have been easily amused by a box of Swan Vestas, but The Prodigy haven't moved into the world of open source software with their arsonistic antics. This *Firestarter* assists with security, a never-ending job. With crackers and 'script kiddies' running amok on the Internet, and vulnerabilities discovered daily as new software is released, keeping a box tightly secure requires forethought and constant attention. Firewall tools are one aspect of this, and we've looked at a couple before in Hot Picks; reader Julian Rigby asked us to give *Firestarter* some coverage, and a fine suggestion it was too.

As it's built around GNOME 2, you'll need the relevant development libraries and headers to compile

from source. That shouldn't throw up any difficulties, but we've also included RPMs for Red Hat and Mandrake on the CD if you'd rather go the binary route. An older version is also available for GNOME 1.4 installations on the project's website, and both will work with kernels 2.2 and 2.4.

You'll need to launch *Firestarter* as root, and although graphical Wizards tend to be sneered at for adding a glossy front-end to inherently complicated concepts, the design here works well. Each stage of Firewall configuration is neatly organised with appropriate info text – it's not quite as helpful as some of the distro-supplied tools, but flexibility comes as a result. After choosing the services to be available and ToS filtering (for giving certain types of service higher priority



Far from the "trouble starter" of the song, this'll keep you out of trouble.

etc.), ICMP filtering can also be tuned and the Wizard finishes.

In normal use, *Firestarter's* main window displays a constantly updated list of port requests; helpfully, each column can be sorted, and it's an excellent way to identify malicious activity. The Prefs box allows settings created in the Wizard to be changed,

in addition to extra goodies including NAT (for sharing connections), while the well-written online manual is suitably thorough. There's no doubt that experienced sysadmins may prefer to craft firewall rules via the command line, but for a quick, simple and very friendly solution, *Firestarter* is top-notch.

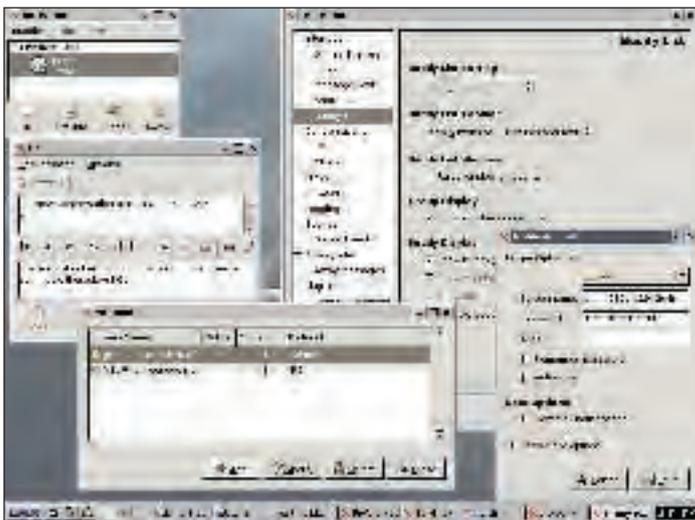
INSTANT MESSAGING CLIENT

GAIM

■ VERSION 0.71 ■ WEB <http://gaim.sourceforge.net>

Life's difficult for coders of Open Source instant messaging clients. Many of the popular protocols are guarded tightly by the service companies, and

developers need to laboriously reverse-engineer any new changes so that users can natter away with free software. *GAIM* is definitely one of the



Buddy list, chat box, account settings and preferences window in action.

most successful open source instant messaging (IM) clients – we last looked at it over a year ago in *LXF36's Roundup*, and since then it has seen lots of work and we've had several requests to cover it again.

GAIM uses GTK2 as the toolkit for its front end, but it'll work fine in any desktop environment or window manager. A Win32 version also exists; considering that every major IM service already has software for Windows, the fact that users want a port is a sign of *GAIM's* quality. A bunch of protocol plugins are supplied with the source tarball so nothing else is required, and an RPM for Red Hat 9 can be found on our coverdisc too.

When first started, *GAIM* pops up a simple login window for selecting an account – new accounts can be added through a friendly setup box, and it's all neatly constructed and very easy to use. Supported protocols in

the stock build are MSN, Yahoo!, AIM/ICQ, Jabber, Gadu-Gadu, Zephyr and IRC, covering all of the popular chat services and smaller networks too. At the time of writing, *GAIM 0.71* still supported *MSN Messenger*, but you should check the website to see if this is still the case. Being able to chat using multiple protocols from one single application is enormously useful.

Tabbed windows help to avoid clutter when talking, while the Buddy Pounce feature is a superb addition for performing an action when a co-chatter does something. Text replacement, previous chat history, mouse gestures and other plugins increase the client's versatility too, and since *GAIM* was featured in our Roundup, various improvements have been made to protocol handling, while full KDE/GNOME systray support, new translations and stacks of UI polishes have been added on top.

Right now, *GAIM* is still the best all-rounder for Linux instant messaging (although *Kopete* is right on its heels – see <http://kopete.kde.org/> for more information), and the friendly front-end, variety of options and top-notch stability ensure a great future. Very highly recommended.

PROGRAMMER'S EDITOR

CUTE

■ VERSION 0.2.1 ■ WEB <http://cute.sourceforge.net>

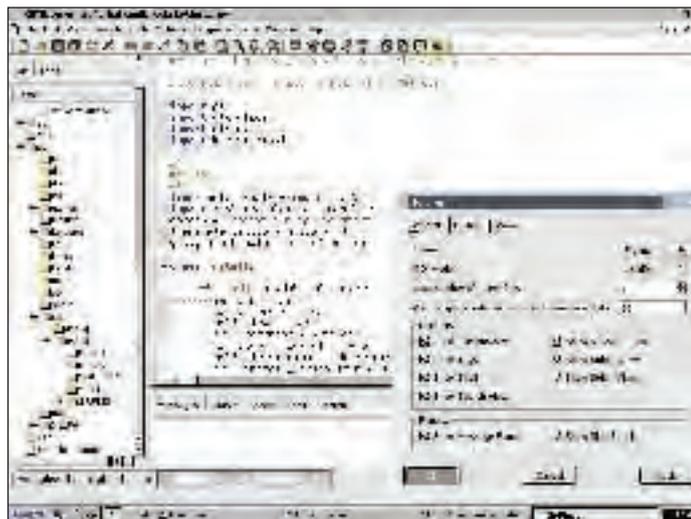
To developers, a text editor is like a home. It's a place in which the user feels comfortable, a place for doing work without pointless distractions, a place where everything feels 'right' and natural. One of the reasons *Vi* (and particularly *Vim*) has remained popular is that it doesn't bombard the user with feedback; equally, *Emacs* lets users create their own working environment. On the graphical side, newer entrants to the editing scene such as *CUTE* are becoming popular, and here we're looking at an early – but fully usable – release of this Qt-based editor.

Version 3.x of the Qt toolkit is required to build and run *CUTE*, and it also uses the acclaimed *Scintilla* source editing component for the meat of the program. Although it's built around Qt, it's not a KDE program and will happily work on other desktop environments. We've also provided a binary RPM on the coverdisc for quick installation on most modern distros.

CUTE's main editing window sports the usual furniture of menu, toolbar and resizable multi-paned coding components – these include a directory tree for picking out files, an output box for compilation messages, and the *Scintilla* editing section. Plus, a taskbar does the job for hacking on several files simultaneously, and on the whole it's a workable setup.

Syntax highlighting for a truly impressive array of languages is supported, with C/C++, HTML, Java, Perl and Python being among the most notable, and they're all tunable in the program itself. *CUTE* makes heavy use of Python for much of its back-end and configuration settings – thankfully, just about everything can be tweaked quite easily without the need for hand-hacking the config files.

CUTE bundles in many editor must-haves, including a macro recorder, unlimited undo/redo, regex find/replace and session restore



CUTE hard at work on a Python file with the Settings dialog popped up in front on the right-hand side.

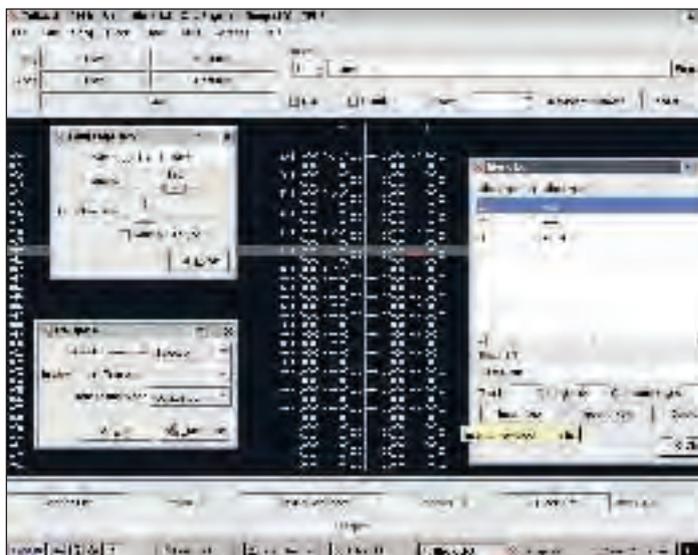
features all topping the impressive list. In our tests, it proved to be stable and speedy, and best of all it's supplied with a healthy range of keyboardings.

CUTE certainly has great potential and naturally fits in well with KDE, although with some themeing it's a solid choice for any coder's desktop.

MIDI SEQUENCER

Tutka

■ VERSION 0.11.0 ■ WEB www.nongnu.org/tutka/



Our rough attempt at a Bach piece, with various control boxes on top.

Way back in issue 14, our *Roundup* looked at a variety of the best and most intriguing sound-related software for Linux. We concluded by saying that the current range of audio tools was satisfying and covered most bases, although many still needed a touch more spit-shine. A few years on, and things are healthier for Linux-using musicians – while the OS doesn't quite rival Windows and Mac OS in this department just yet, especially in the sequencing arena, new tools such as *Tutka* will hopefully rectify this.

Tutka is a MIDI sequencer with a tracker style front end (ie it's similar to OctaMED and SoundTracker in use) and runs under GNOME 2.x. Compiling from source shouldn't require any somersaults providing you have the necessary development bits for GNOME – it's always good to see new software that doesn't rely on a zillion obscure dependencies.

The cleanly laid-out GTK2 interface holds just the bare essentials; it's nothing like as cram-packed as the excellent *SoundTracker*, but sadly the scrollbars aren't functional at present

(so navigation has to be done by hand), nor is there any sane way to change the minuscule default font in the tracker component. The keyboardings are probably insufficient for power-users too.

Songs are built up from blocks of varying length, so repeating sections can be created easily, and transpose, shrink/expand and copy/paste functionality for blocks and tracks has been coded in too – it handles much like a traditional sample-based tracker. Editing, defining instruments and adding MIDI commands is all enjoyably simple, and the native XML file format is just about readable in the flesh.

Tutkaplayer is available for playing songs without firing up the whole app, and planned for future releases are a proper online help system, options to modify the keyboard configuration and ability to work on multiple songs simultaneously. There's not a huge amount of features and it certainly needs the odd bit of work in places, but in our testing we found it to be reliable and fast. If you're a MIDIer, it's worth giving it a test run or two.

KDE MEDIA PLAYER

AmaroK

■ **VERSION** 0.6.91 ■ **WEB** <http://amarok.sourceforge.net>

As Linux marches steadily

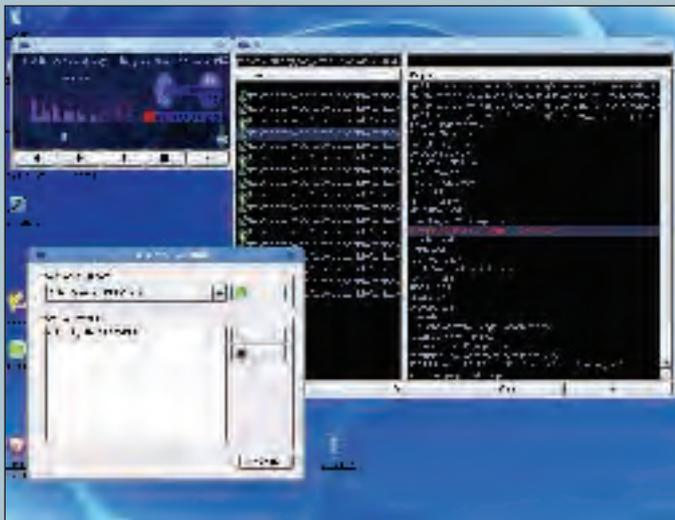
onwards in gaining desktop acceptance, multimedia apps are one area in which the OS is becoming ever stronger. We've seen recently how *MPlayer* has been immensely useful in allowing free software fans to play all manner of file formats, and *XMMMS* is still the most popular music tool around. KDE's army of coders haven't stood still, though – we were pleased with *JuK* in issue 41's *Hot Picks*, and a more recent entrant to the audio scene is *AmaroK* (the final letter is always capitalised).

Currently only available in source form, compiling *AmaroK* shouldn't pose any difficulties providing you have *Qt 3* and *KDE 3.1* installed and their relevant development bits and bobs (*kdelibs-devel* etc in many

distros). Note that the standard `./configure, make` etc process in the base directory only builds *taglib* due to a glitch; you'll need to `cd` into *amarok* and do `make` and `make install` again.

AmaroK's ultra-slick interface uses dark blues and a smooth scrolling song title display, with a mini volume control slider and quick access to the playlist (although the equaliser code isn't in place yet). It may not rely on the standard KDE file selector component, but *AmaroK*'s playlist editor allows collections to be built up with drag-and-drop, and the quick-search box at the top is a handy little extra.

With *aRts* as the back-end, *AmaroK* will play MP3, Ogg Vorbis and audio CDs, and stores its playlist in the popular M3U format in `~/.kde/share/apps/amarok` (.pls will work too). Besides the basic list



A smooth AmaroK setup showing the main window, playlist editor and effects box.

managing and tune-playing functionality, the application also boasts a crossfade mode (for seamlessly moving between songs) and a sprinkling of effects thrown in for good measure. Pleasingly, keyboard shortcuts can be redefined through the program itself – always good to save faffing around with the

mouse – but there's little else in the way of configuration.

That's not a problem though, as *AmaroK* just strives to be a simple and easy music player with easy-going playlist management features. It's not as fully-fledged as *XMMMS*, but fits in very comfortably with KDE. Organise those .oggs today!

NETWORK TESTING TOOLBOX

Netwox

■ **VERSION** 5.2.0 ■ **WEB** www.laurentconstantin.com/en/netw/

All experienced admins have a

suite of favourite trusted tools for monitoring their network and making sure everything's hunky-dory. Linux systems are typically supplied with a decent range of utilities, but reader Darren Aspin asked us to give *Netwox* a quick run-through as an alternative. It's described as a 'network testing toolbox' and runs under UNIX flavours and Windows.

In order to get the full suite up and running, a bit of planning is required – you'll need the three main tarballs (*netwib*, *netwox* and *netwag*) at hand and you'll have to install *netwib*, the library component, first. The only major dependency this requires is *netlib* as provided on our CD, and Tcl/Tk 8 is also needed for the graphical front-end (this should be supplied in most distros). Thankfully, the non-standard build procedure for

each package is well documented and it doesn't take long to get everything installed correctly.

In structural terms, the *netwib* library provides facilities for converting addresses, manipulating packets and traffic sniffing etc, and the *netwox* program links with this lib to make a usable little tool. Lastly, *netwib* sits on top as a graphical front-end. General use of the *netwox* binary is with a command number and associated parameters; just running the program on its own throws up a mini help system for locating the command number needed. For instance, **netwox 1** prints out the local system's network config, while **netwox -3 <IP>** shows hostname and MAC address details.

With 101 built-in commands, it's extraordinarily versatile and should satisfy all but the most demanding of network admins. Alongside the sniffing,

packet fiddling and information retrieval tools, mini versions of *ping* and *traceroute* are also available. Some users may criticise the fact that there are a few commands with potentially malicious purposes – at the same time, though, as Dan Haagman of *TSafe* explained in *LXF46's Linux Pro*, these can be supremely useful in finding weaknesses in a security

system, so it would be fair to say that they're not always put to bad use.

Netwox may not be the most accessible tool ever, but the documentation explains many common problems, and considering the small size of all three components, *Netwox* (and its *netwag* GUI) are certainly worth having around if you manage a network.



Netwox, Netwox's help screen and a sniffing operation on a network.

TETRIS CLONE

LTris■ **VERSION** 1.0.5■ **WEB** <http://lgames.sourceforge.net/index.php?project=LTris>

Alexey Pajitnov's *Tetris* is one of the most important games in computing history. Simple to grasp, life-destroyingly addictive and immensely challenging on the later stages, it was almost entirely responsible for Nintendo's overwhelming success with its original Game Boy – the ultimate 'killer app' as many have labelled it. By transcending genres and being attractive to players of all ages, *Tetris* has seen an endless number of clones (some such as *TetrisSphere* on N64 took the idea too far) and Linux variants abound. *LTris* is one such clone and part of the *LGames* suite.

You'll need SDL to compile *LTris* from source – that's installed by default on all major distros so no difficulties there (and if you need a newer version it can be found in the Essentials/SDL/ directory on our

coverdisc). The RPM we've included has been built for Red Hat 9, and with the very few dependencies it should work fine on other recent distros too.

Presentation wise, *LTris* isn't particularly engaging; still, overloading a relaxed puzzler with rendered intros and pointless frills is in most cases a shooting offence. The menus offer help bubbles to explain all the options, and block previews and guides can be enabled for inexperienced *Tetrisers* (with an appropriate penalty on scoring). Sharp images for the blocks and suitably clunky sound effects all work well in play, but it's a shame there's no chilled-out soundtrack drifting around in the background.

Alongside the correctly implemented 'Classic' *Tetris* mode, a few multiplayer game styles are available (humans and/or CPU) and an intriguing 'Figure' option in which levels begin with a block-



Your Hot Picks man mike being soundly thrashed by the CPU's rocketing speed here. Oh, the humanity!

built picture in place (car, castle etc) It adds a fresh challenge on top of the game foundations without ruining what makes it great – a mistake so many software houses have made in the past.

LTris's CPU opposition works far too fast to make it entertaining, but the Classic game hasn't been tampered with and the Figure mode is a sweet addition. All in all, it's good

old *Tetris* and that's what matters – still horribly addictive yet curiously soothing. Just don't follow the example set by Faiz Chopdat of Lancashire, who was jailed for four months after refusing to stop playing *Tetris* on his mobile phone on a flight returning from his honeymoon in Egypt. The date of his infraction? *September 10th, 2001...*

SCORCHED EARTH / WORMS TYPE GAME

Atomic Tanks■ **VERSION** 1.0.0-RC2 ■ **WEB** <http://atanks.sourceforge.net/>

Human nature tends to involve a lot of this: people being happy at

someone else's misfortune. Watching a chap slip on a banana skin, giggling as a



Pah! I was just about to unleash my world-shatteringly powerful nukes too.

sworn enemy falls down a manhole when you're being chased, or laughing nastily as a cycling race competitor get a broom handle stuck between his spokes – the amount of endorphins released in such moments is incalculable. Those clever Germans have even thought up a word to precisely express this feeling – 'schadenfreude'. Because of this, multiplayer games are very good at improving contentment and establishing the pecking order, and few do the job as well as *Worms*.

Atomic Tanks is an implementation of the *Scorched Earth* and *Worms* series of games; summed up, you're positioned on an undulating landscape and have to blast your opponents into oblivion. The challenge comes through judging the trajectory of your projectiles – annoyances like wind and gravity get in the way of your planned violence, so careful timing and accuracy is the order of the day. *Atomic Tanks* needs the Allegro library to build, and many distros have this in their archives, but otherwise building from the coverdisc's source tarball should be no hassle.

Before each bout, you're offered the chance to buy new weapons with your cash reserves (cycled through in-game with the Tab key). These range from small popping explosives to thunderous megabombs, with shields and teleporters also at hand. Weather plays a big role in the proceedings – thunderstorms and meteor showers can be enabled, and as the landscape crumbles quick action is essential. Visually it's no great beauty and the sound effects are as you'd expect, but it's the well-crafted gameplay that shines through.

CPU players make entertaining quips when suffering a hit: 'Your father was a hamster' is what anybody would say during a battle. Unfortunately, there's seemingly no way to resize the in-game font, and it's painfully small at higher resolutions. A network mode is missing too, but luckily, that's on the cards for a later release. Overall, *Atomic Tanks* is top-notch projectile-based fun, especially with many human competitors, and once network play is added and some extra polish applied, it'll be marvellous. **LXF**

Free your office with... OpenOffice.org

Many users think that that *OpenOffice.org* is the most important Open Source application in the world at the moment. "Quite probably," says **Richard Cobbett**.



While the exact truth of the above statement does depend largely on your own personal requirements – *Apache* developers are unlikely to be quite as excited – it's certainly true that there are few other apps which command quite such wide-ranging impact. Linux itself? Amazingly important – we all know this, but it's *OpenOffice.org* which is more likely to turn heads of those in the average office without the will or drive to switch their entire OS. That's a serious investment in both time and resources, while *OpenOffice.org* can not only be downloaded and tested on any Windows system before taking the plunge, it offers many compelling reasons to do so – from philosophical

freedom down to free availability. *The GIMP*? Certainly a powerful package, and likewise available cross-platform, but the learning curve and interface put it on the back-foot, while you could happily sit your boss down in front of *OpenOffice.org* and have him tapping away by the end of the afternoon. There are no obscure commands to remember, no unwieldy interfaces that try to turn basic functions into *The Krypton Factor*, no cute developer quirks to put them off and no hint that this is anything other than a professional level suite for serious users. And of course, it isn't cursed with a name that will have him or her worriedly exclaiming "You want me to try *what?*"

Four-pronged attack

Put simply, *OpenOffice.org* has one simple goal: to oust Microsoft from the desktop. Its trusty weapons are a comprehensive suite of four key applications: the word processor *Writer*, spreadsheet *Calc*, *MS PowerPoint* replacement *Impress* and vector-graphics package *Draw*. Notably missing are an email client and dedicated database (*StarOffice* does ship with a dedicated database component, but still no standalone application), but these are hardly in short supply on Linux. The whole lot is downloadable from www.openoffice.org, although versions ship with almost every major distribution. If you are still without, have no fear – the installation is a picture of how all software should get onto your PC, happily installing on almost anything. *StarOffice*, currently at Version 7, is the commercial release – or, to be more exact – a commercial release. Here's where things get a mite complicated: Sun's *StarOffice* suite has been around for many years, gaining some attention, but little to give Bill any sleepless nights. A couple of years ago, Sun decided to open-source the suite, creating the *OpenOffice.org* project that we see before us. Rewritten and upgraded, the free *OpenOffice.org* is now the base of the commercial *StarOffice* – along with a handful of extra goodies to sweeten the deal. Sun retains the trademark to this for obvious reasons, but there is absolutely nothing to stop you from grabbing the code and spinning up your own alternative versions should you so choose. Ximian took advantage of this for its *Ximian Desktop* product, ripping out and replacing the standard graphics

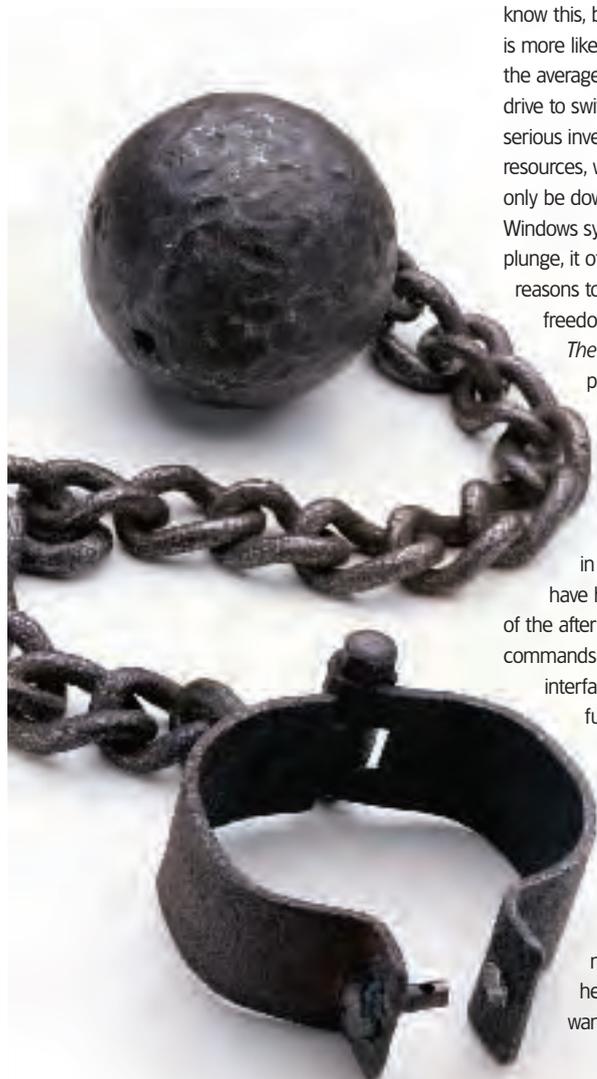
to fit its customised look. Others simply use the base code, but alter its name to "Word Processor" on the K menu. As with all good open source software, it's a simple matter of choice.

No lock-in

This is of course one of *OpenOffice.org*'s key selling points. While *Microsoft Office* does its level best to lock you into individual systems, *OpenOffice.org* advocates open standards at every point. You can see this for yourself – save a document in its favoured .sxd format, rename it to a .zip and open it up in any compression utility. You'll find all the pieces of your document lying in front of you in standard XML format, not obfuscated to hell and back to ensure that only *OpenOffice.org* can read it successfully.

Meanwhile, the suite will happily open up almost any basic *Microsoft Office* files – an average document should be seamless, although more complicated files aren't guaranteed to come through perfectly. Writing to these formats is of course much more reliable, and in over a year of using the suite professionally, we have never encountered any problems switching files between our Linux machines and the Windows systems in the office. In at least one respect, we actually have a major advantage: *OpenOffice.org* exports straight into machine-independent PDF format, while even the just released *Microsoft Office 2003* still insists that everybody uses .DOC.

While the suite does have a fair way to go before achieving perfection, we'll be seeing over the next few pages just how powerful it really is. If



you've been looking for that Linux killer application, look no further – *OpenOffice.org* is here, and here to stay – and not just for Linux, either!

WRITER

Writer is a word processor. At a basic level, that's all you really need to know. You have fonts above, text below, and reams of menus to convert a simple page of text into a professionally laid-out document. Its day to day features will offer little to surprise you – but its many little eccentricities make for a unique experience. The first that you'll discover, possibly with a touch of frustration, is *Writer's* love of auto-completion. As you type a word, *Writer* attempts to work out exactly what it is that you're typing, and fill in the remaining letters from its dictionary. If this is correct, you hit

Return to snap the word into place, if not, you simply continue typing as normal.

This is a splendid time-saver when you get used to it, but as close to hypnosis as most touch-typers will ever reach – blazing fingers rattling off the wrong word on reflex. This duly switched off – for a while, at least – but little in the rest of the package will throw you for long. The one key omission, which we frankly cannot believe is still absent, is a word count. Hunt around and you will find some basic information in your document's Properties, tucked

away just behind the plans to bulldoze Arthur Dent's house, but even this only works on entire documents and not simple selections of text. Until this is fixed, you have to copy any text that you wish to count into a separate document and word-count it there – a complete pain in the underwear.

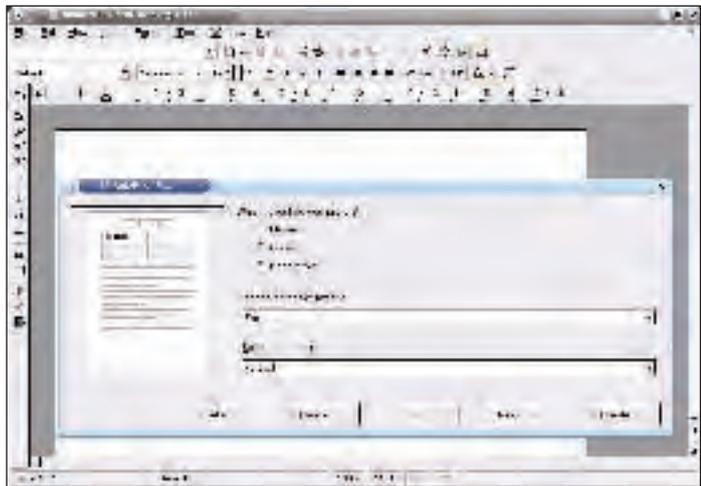
Joy to use

Luckily, it is the only one that has plagued us in over a year of day-to-day use. In every other respect, *Writer* is a joy to use – and much more powerful than it initially appears. For starters, in addition to standard features such as Mail Merge, it gives you the ability to include specific fields and even conditional documents – only showing text if the user needs to see it, and performing

calculations live in your files, be it mathematical sums, or altering 'page' for 'pages' as your document grows. Likewise, you can group

multiple documents into a single Master file, ideal for either long documents like books, or projects with multiple authors, and track any changes made to them.

OpenOffice.org has an unfortunate tendency to hide its light under a bushel in these cases, and it's only when you look at the help-file's comparison of terms between it and Microsoft Office that you start to see just how much you've missed – be it



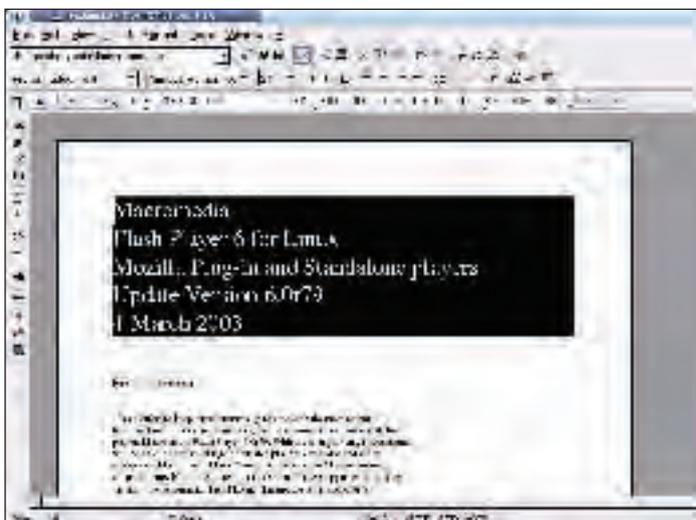
password protecting individual sections of a document or realising that *OpenOffice.org* calls its pronunciation guide system 'Ruby'. While you may not be able to find a direct replacement for every single feature that you use in Microsoft's behemoth, you will find the vast majority present and correct.

CALC

Second only to *Writer* in the importance stakes, *Calc* offers everything that you need in a spreadsheet – complete mathematical processing, neatly laid out across

Document templates are provided in the form of AutoPilots. There aren't many of them, but they come in extremely useful when starting out. The spelling checker quickly learns your favourite words, underlining mistakes as you make them. The Gallery contains plenty of sample imagery, including these basic 3D effects to drag and drop into position.

'If you've been looking for a Linux killer app, look no further – *OpenOffice.org* is here, and not just for Linux either!'



Maintain consistency across a document using predefined font settings and the floating Styles panel.

multiple pages on the standard 32000 lines. While its basic subject matter is rather more intimidating than the other packages on offer, *MS Excel* users will be able to jump straight in – the help files quickly bringing everybody else up to speed on both the basic methods of laying out a spreadsheet, and the more in-depth options on offer, such as assigning conditions to colour-code key cells (most obviously, negative numbers lighting up a bright red). Each individual cell brings a raft of new options, including basic font and colours, the ability to rotate text to any angle and a choice of data formats (such as Currency and Date, complete with the correct English form) along



FREE YOUR OFFICE



with more advanced features, such as cell protection, hiding the formula and not printing out the result when a hard-copy is produced, unless the user punches in the password first.

Data mining

Naturally, graphics are a key part of any spreadsheet, and *Calc* certainly doesn't disappoint, bringing all the usual chart types from 3D pie-charts to the polygonal lines of a net

– even if the finished results do look a touch rough-and-ready.

Data mining is also on offer in the form of *DataPilot* – producing

interactive tables to zero in on the facts and figures that you genuinely need, instead of having to thumb through page after page of data. This is extremely easy to use, even just following the basic information in the online help files, using standard conditions to filter through your spreadsheet and making removing duplicates or case-sensitivity as easy as toggling the appropriate option.

Actual data entry can be done the old fashioned way, simply punching values into the right boxes, or by directly importing your existing *Excel* files with *OpenOffice.org*'s built in filters. Another alternative is to make use of external data, pulling in tabular information from regular HTML pages. These can either be static pages on your main machine, or live web pages – with the option of updating their content second-by-second if need be, or at any other interval you choose.

While *Calc* has – in our experience

– proven to be one of the more finnickily parts of the *OpenOffice.org* suite when it comes to directly importing *Excel* data, it certainly brings enough power to the table to make it worth trying out your documents, or moving them across completely.

IMPRESS

It's best to keep your eyes glued shut the first time that you use *Impress*, because the sample templates that it comes with certainly won't. If it falls at any hurdle, it's not that it's in any way incapable of producing great looking *PowerPoint*-style presentations, but rather that *PowerPoint*'s success largely stems from its total ease of use. Slap down a nice looking background, dig out the clip-art, scatter a few buzzwords around the place and even the most technophobic executive is ready to torture their captive audience. *Impress* makes you work a little harder, and feels considerably older than the rest of the suite – and not, sadly, in the good old 'tried and tested' way. For starters, the interface reaches all-new levels of clunkiness, forcing you to skip between entire workspaces just to bring up something as simple as a thumbnail view of your slides, and trapping the playback controls on the horizontal scrollbar. Edit a graphical object's style, and a mammoth tab collection blasts out of the screen. Serious refinement is needed across the board to make *Impress* as comfortable to use as its Windows-based competition.

Get past this however, and you'll find it packed with excellent features. In addition to the standard import/export of Microsoft files – all the more important, considering that 'PowerPoint' has become the generic term for computer-based slideshows – it mirrors its proprietary rival almost perfectly, with a huge range of page styles to pick and choose from, all of the standard effects and tools to handle the parts of the presentation that your audience never sees – such as timing your reading and creating custom slideshows by picking and choosing groups of slides for each audience's perusal. The finished show, in addition to being played back in either *PowerPoint* or *Impress*, can then be saved as a Flash format .swf,

again ensuring compatibility across the board. And that board's not just the Board of Directors.

Impress's results are perfectly acceptable, and are almost indistinguishable from the 'real thing'. While it offers few compelling reasons to switch across if you already have an equivalent tool on hand, it is an essential presence in the overall suite, and could quite easily rise in the ranks to essential with a few better samples to play with, and an overhauled UI to ease presentation construction.

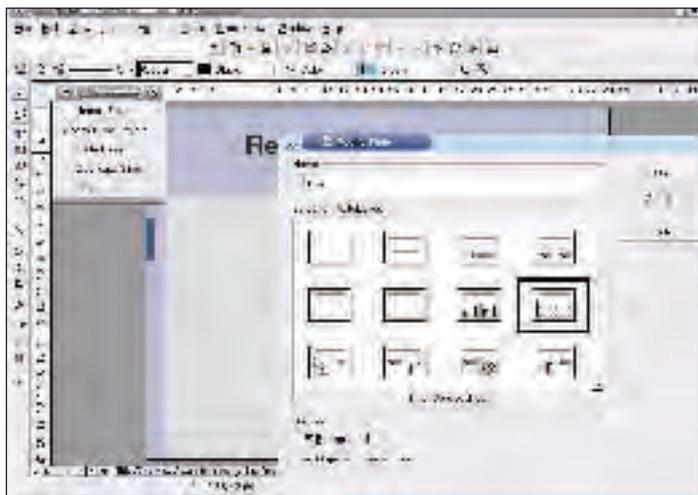
DRAW

Draw often ends up being passed over when the *OpenOffice.org* suite is discussed, if not through any real failing of its own. At the time of writing, even the official *OpenOffice.org* forums continued this trend – *Draw*'s corner of the web carrying just 63 topics to *Writer*'s 930. Ultimately, all that can really be said about it is that it's a competent enough vector drawing package. Competent enough. Those words are a killer. It will never find itself being held up as a serious alternative to *Freehand* or *Illustrator* – unlike *The GIMP* with its *Photoshop* comparisons – nor does it really stand up against other established packages like *Corel Draw*. It's simply an okay program for creating simple vector-based artwork, with a couple of handy features like connectors for putting together organisation charts, and the ability to render (very) simple 3D objects such as spheres and cones, and convert text objects into models that can be lit and rotated in aesthetically suitable ways. And that really is more or less it for this part of the suite.

Quick fix

If *Draw* lacks any exciting features, it does at least bring most of the standard ones to the table. Cross-fading offers you an easy way to morph one shape into another, with gradient fills helping to break up the typical flatness of vector based imagery. *Draw* is also likely to be your first real introduction to the Gallery, which stores all of *OpenOffice.org*'s supplied clipart and other media files. It's also worth noting that several of its drawing tools are available in *Impress* for live-artwork in your presentations,

Impress offers few proper templates, but there are plenty of slide layouts to start working on. Animation, graphics and sound can all be combined into Impress slideshows so good looking that they could almost be PowerPoint.



but sadly the ability to use *Draw* without actually opening it doesn't offer a terribly compelling reason to do so. In the end, all that can really be said about it is that it's there if you need it, which won't be much longer than it takes to download a dedicated vector artwork package. In the meantime, artistically minded users will at least be able to use it to bang out quick shapes to jazz up their documents, if probably not create the Linux world's Mona Lisa – especially with the likes of *Sodipodi* readily available. Perhaps this is its biggest catch: it's simply not easy enough to use if you aren't already a decent artist, but it offers nothing to inspire you to change track.

MISCELLANEOUS FEATURES

Beyond the basic software packages on offer, *OpenOffice.org* is bursting with powerful features to make it easy to use. Most important of the lot, unsurprisingly, are the *Microsoft Office* filters on offer. Like it or not, these file formats – bloated and twisted as they are – have become the standard format for business use, and this is often where competing suites fall on their faces. Luckily, *OpenOffice.org* is not only capable of reading in almost anything from an *Excel* spreadsheet to a *Word* document, it does so extremely well. We know several people who have even tried shovelling a mangled file through its filters, claiming that *OOo* does a better job of pulling out any remaining content than *Office* itself. Of course, the conversion will only go so far – plain text and layout typically comes through almost flawlessly, but more complicated documents and those relying on Microsoft technology like SmartTags and *Office* macros can still throw a monkeywrench into the works. On the other side of the argument however, we have never experienced any problems loading a saved *OpenOffice.org* .DOC file into our word processors, and we've given it plenty of time to throw a fit over the years.

Actually creating a new file is as simple as opening its part of the suite, be it *Writer* for text documents or *Impress* for *PowerPoint* presentations, but once inside the suite, it's

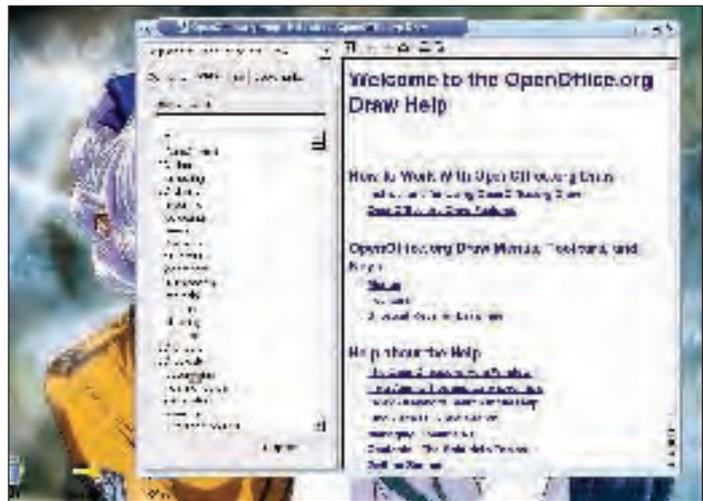
extremely easy to chop and change between applications. Templates are provided in the form of AutoPilots – Wizard driven interfaces that set up basic pages, and basic is the key word. Don't expect anything jaw-dropping to emerge from any of them, but there's little to complain about if simply trying to create a quick letter, web page, or even Euro-based currency converter.

Options

The Options menu provides several ways to ensure that you don't encounter any problems, including activating auto-save every n minutes, creating backup copies and picking default file-formats. Open or not, *OpenOffice.org* is fond of pushing its own favoured formats – including a useful, if occasionally tooth-grinding dialog box on quitting to remind you that saving in external formats may have resulted in data-loss. As already mentioned, any documents can also be separately exported as PDFs, removing any such risk and ensuring that people will be able to read your documents regardless of platform (including on Mac OS, to which *OpenOffice.org* has very recently been ported across).

If you can't use all of *MS Office*'s functions, you can at least use the *OpenOffice.org* replacements. Macros, for instance, are never more than two clicks away, with an option in the Tools menu to start recording, and a Stop button to tap when finished. If this is too high-level for what you need, and yet you have no desire to dive into the source itself, there is plenty of middle-ground in the *OpenOffice.org* API/SDK, coding additional features in C++, Java or the remarkably familiar *OpenOffice.org* BASIC. This has its own dedicated site, which you can find at <http://api.openoffice.org>, although you will need a fair amount of programming experience to make good use of its features.

Keeping track of addresses and contacts is extremely easy, courtesy of the ability to import *Mozilla/Netscape* Address Books – or hook the suite up to an LDAP source. *Mozilla* is generally good at importing these from other third-party clients, saving a lot of tedious data-entry. Naturally, your own information is much easier to access – tucked away in the User



Information panel. This is also where you can customise the individual applications on your machine, deciding whether to open up *Evolution* or *KMail* to send an email, or choose between *Konqueror* and *Firebird* when you click on a hyperlink. Naturally, all of these settings are valid across the entire suite, along with important built-in elements such as the dictionary and thesaurus. The latter is however slightly clunky to use, lacking the ability to simply punch in a word (instead forcing you to select one and then load the thesaurus) and provide a dictionary definition of exactly what you're about to offer up as a replacement. With even a word as simple as 'test' throwing up choices ranging from 'catechize' to 'joust', a little button marked 'Huh?' to link to Dictionary.com or a related site

The online helpfiles are exceptionally good, although occasionally they could point towards more exact destinations.

OpenOffice.org BASIC has its own dedicated website, although you'll need coding skills to make good use of it."

wouldn't go amiss. We're certainly curious to know why 'chicken' would offer up 'floppy', 'poltroon' and 'Brahma' instead of pushing the more conventional 'afraid', 'scared' and – our favourite – 'milquetoast'. Naturally, *OpenOffice.org* isn't aimed solely at the English-speaking world. In addition to translation packs for these core features, and 23 versions available, the suite has full Unicode support for any additional languages that you may need to plug in.



« StarOffice 7.0

While Richard Cobbett is always loathe to pay cash for software, StarOffice makes a very tempting play for the contents of his piggy bank.

When looking at *StarOffice*, it's hard to get past two key points. Firstly, almost everything it does is already available for free. Secondly, one of *StarOffice*'s selling points is that it's not available for free. Bizarre logic it may be, but there is a perfectly valid reason for it – even discounting issues like support and extra features, plenty of people out there automatically associate 'free' with 'cheap and nasty'. Putting a value on *StarOffice* immediately gives it credibility as a commercial package, making it a much easier sell to a company's accountancy team. It's important to note however that *StarOffice* is now based on *OpenOffice.org*, not the other way around – should Sun decide to shelve the commercial release, there is nothing preventing its free version continuing for many years to come. Because of this, while *StarOffice* itself may be limited to a 30-day trial version, you can easily test out most of the new features from the *StarOffice* 6.x line before taking the plunge. Equally, while Sun is unlikely to thank you for pointing it out, this is largely *StarOffice* 6.2 –

original plans to wait for the next big release of the suite (codenamed 'Q') having been put on ice in the name of a milestone release.

As with *OpenOffice.org*, *StarOffice* consists of three key apps: *Writer*, *Calc* and *Impress*, with *Draw* bringing up the rear. You'll notice the continued absence of a database app, though some basic functionality is provided via the *Adabas D* component. This is not hugely relevant however, as a Personal Edition (limited to three

users) is downloadable from the Adabas pages that are found at www.softwareag.com. The other features on offer, while useful if you need them, remain firmly in their niche market – a *WordPerfect* filter to back up existing *MS Office* converters, and a handful of tools for handling Asian fonts and docs. Add a sprinkle of templates, a replacement spellchecker and thesaurus, and that's the end of the real extras.

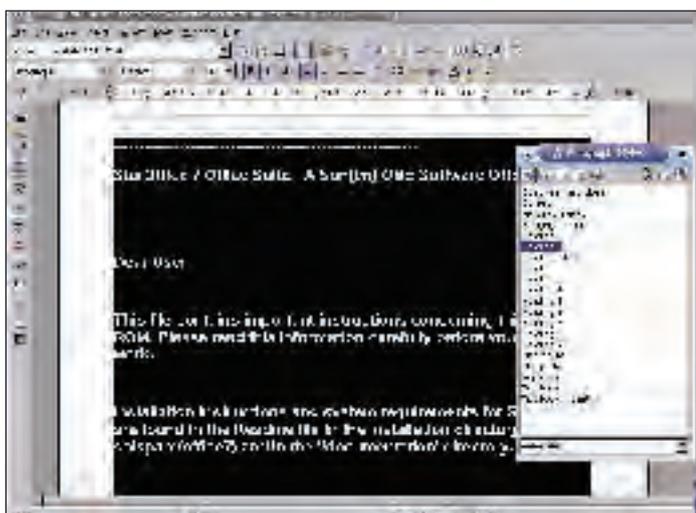
Extra features

This said, *StarOffice* does come with one key feature that *OpenOffice.org* lacks – support. In addition to the live Help file docs in the main package, you get a standard dead-tree manual to plough through, 24x7 web based support and access to a main helpdesk to get you through any early stumbling blocks. Enterprise customers can take a step up from this to full level support, although of course this costs extra. It's worth noting that both *OpenOffice.org* and *StarOffice* share an extremely easy installer, making it straightforward to get up and running on either Linux or Windows platforms, meaning that sysadmin or user alike will be able to get started with few headaches.

To recap, *Writer* handles almost any *MS Word* requirements with aplomb, although DTP isn't its forte (especially if current plans to merge the dedicated *Scribus* with the *OpenOffice.org* suite



All installation programs should be this easy. *StarOffice* really raises the bar for other commercial software packages.



Spot the difference? Beyond the words '*StarOffice*', there is little to visually separate the two versions of Linux's favourite office suite.



come to fruition). Word-counts remain the most notable omission, and certainly top on our list of elements that need to be fixed ASAP, but the rest has more than proved itself as a stable system, with all of the features that you'll need for day-to-day use, be they scripting, mail merge, or even groupworking. Many of the features are hidden out of view, or obscured by simple dint of not sharing the standard Microsoft names, but a look-up chart in the Help files quickly points you in the right direction. *Writer* can refresh pages from the Internet, track changes between docs and keep track of different versions via a splendid Properties window. Beyond this, the most obvious unique feature is the floating Styles window, which makes it easy to maintain consistency even across the longest of documents (themselves potentially split up and merged together as a Master Document). Once a document is finished, it can be exported in any form, from *Microsoft Word* to Adobe PDF, or *StarOffice's* own XML-based format. *Writer* is an excellent package, whether free or commercial.

Calc in the downloaded retail version of *StarOffice* doesn't supply the manual, but its familiarity and the ability to pull info in from online sources and conduct basic data-mining make it more than capable of juggling any data that you throw its way. It is best used from scratch, building up spreadsheets specifically for use within *Calc* rather than importing them from *Excel* – though this is purely a matter of precaution, and the import filters are powerful enough for most general users.

Impress is your handy, ever-present *PowerPoint* package. It may not be the

TIPS

- 1** You can insert calculations into any document. Simply point where you want them, press **F2** and enter the necessary formulae. These are in exactly the same format as their dedicated cousins in *Calc*, although of course without having to sit in a cell.
- 2** *OpenOffice.org* features seemingly hundreds of shortcut keys. Look them up in the Help file to speed up your work, especially given that finding many options can involve jumping through many tabbed dialog box-shaped hoops.
- 3** The Help Agent flashes up whenever it has information to impart, but quickly becomes annoying. Turn it off to enjoy a smoother session, as it rarely provides anything of much relevance. At least it's

not the annoying evil paperclip. Other layout hints are the same as under Windows, notably the squiggly red underline of a spelling mistake.

4 *OpenOffice.org* automatically imports any VBScript embedded into a document, saving it out again when you're finished. In addition to this, you can write your own scripts using *OpenOffice.org's* own form of BASIC.

5 Have a lot of documents to deal with? An AutoPilot is on hand to convert everything in a folder into nice, open standard forms. Simply point at the appropriate location and leave it running.

6 Trying to find a particular option? Look for the help topic 'Comparing

Microsoft Office and *OpenOffice.org* Features' in the main documentation – it lists both sides' preferred buzzwords in a simple, easy-to-follow manner.

7 Save time by automatically saving in your favourite file format. The relevant switches can be found in Options> Load/Save, along with a wide range of other settings to filter document types and directories.

8 Need to work in an unsupported language? Find links and advice on localising and internationalising your copy of *OpenOffice.org* at <http://110n.openoffice.org>. Language packs are available for most common languages, although we've yet to see the official Klingon variant.

most streamlined package in the world, but can produce extremely impressive results. As well as the standard *Office* import filters, anything that you create can be outputted into Macromedia Flash format for the web, complete with animation and basic interactivity, and the results are all but indistinguishable from the 'real' thing. In keeping with this theme, *Impress* is the part of the suite that benefits most from its *StarOffice* home, if only for the increased collection of clip-art on the disc, and accordingly, the less time that you need to spend around *Draw*. *Draw* itself remains firmly the black sheep of the flock, with very few genuinely crucial features, although that doesn't mean that it's unusable. With a full range of drawing tools, gradient fills and other In particular, having its connector objects on hand can greatly ease the pain of creating

charts and quick pointers, and the presence of a few of its drawing tools live in *Impress* itself can come in very handy for quick additions to presentations.

In summary...

The big \$640 question is simple: are *StarOffice*, and *OpenOffice.org*, ready for your home desktop and business alike? Yes. While a feature-by-feature comparison of it next to Microsoft's opus may leave it dragging a bit behind, in practical terms, all of the features that you genuinely need have been brought along for the ride – evil paperclips, SmartTags and all the twiddly features that only tiny numbers of folks actually use put aside for the moment. The package as a whole is exceptionally stable, and more than capable, and the price is more than fair, whether you end up paying it or not. **LXF**

LINUX FORMAT VERDICT

FEATURES	9/10
PERFORMANCE	8/10
EASE OF USE	8/10
VALUE FOR MONEY	8/10

The best Linux suite out there, even if you'll almost certainly find that its free version does all you need and more.

RATING **8/10**



ALTERNATIVE OFFICE SUITES

Linux, as always, offers users a choice...

There are plenty of office suites to choose from under Linux, each offering their own unique features and specialities. Most famous is the *KOffice* suite, and few can match its impressive scope: the word processor *KWord*, spreadsheet *KSpread*, *PowerPoint*-like *KPresenter*, flowcharter *Kivio*, vector drawing application *Karbon*, raster editor *Krita*, report generator *Kugar*, charting package *KChart*, formula editor *KFormula* and database tool *Kexi*. All of these are available as free downloads, with others, such as project manager *KPlato*, planned for later release. Like *OpenOffice.org*, this tends to ship with the bigger distributions, although its overall level of quality is not as high. You certainly don't get the same ease of mind when handling Microsoft formats.

More specialist applications are also available in this category, including *AbiWord*. Working purely as a word

processing tool, it does a far better job of matching the traditional *Microsoft Office* look and feel. Capable of importing its rival's documents from the off, *AbiWord* offers strong plug-in architecture to click your most needed features into place, from a thesaurus to a text summariser and spelling checker with support for thirty languages. A full list of these extras can be found at www.abisource.com/twiki/bin/view/Abiword/PluginMatrix. Like *OpenOffice.org*, *Abiword* has the benefit of being available across all three main platforms – Linux, Windows and Mac.

Finally of course, there's the big MS itself. While Hell will be an ice-rink long before an official release of *MS Office* under Linux, it is possible to run the old 2000-era packages with the help of *CrossOver Office* (available standalone from www.codeweavers.com, or bundled

with a number of distributions, such as SUSE Office Desktop). With *OpenOffice.org* at its current level, there is much less practical need to do this these days, but it remains an option if you absolutely have to use the original Microsoft software. Sadly, more up to date versions, including the brand new *Office 2003*, are not supported – yet – but there are still ways and methods, including outright running Windows itself via *VMWare* (www.vmware.com). With Microsoft's continued move towards an *Office System* rather than individual *Office* applications, this may become an important option in the months to come – especially if you need to communicate with companies that have invested in it – but an extremely expensive way of working unless you have a demonstrable need for particular features that it brings to the table.

PREVIEW KDE 3.2

KDE developer **Chris Howells** takes a look at what is in store for the next release of this popular Linux desktop environment.

KDE is one of the largest Open Source software projects around today, and currently its developers are busy working on the release of the next major version.

KDE was originally announced in 1996 when Matthias Ettrich posted to several USENET news groups about a project he was interested in working on, the 'Kool Desktop Environment', a pun on CDE (in Ettrich's native German, 'C' and 'K' are pronounced similarly). Matthias originally asked for around 20 volunteers and very shortly had that number, and so work began.

Another Matthias, Matthias Kalle Dalheimer, an early developer on the project, got the coding started by writing two C++ classes, called *KConfig* and *KApplication*. Both of these classes have been extended since then but are still in use today. It is interesting to note how similar today's *KConfig* class is to that originally written in 1996, a testament to the excellent flexibility that the KDE development platform provides, largely as a result of the use of *Qt*.

Qt is the C++ library on which KDE is based; it is developed by Norwegian software company TrollTech. As well as providing the user interface widgets such as buttons, it also provides a lot more besides, such as data handling facilities. Originally *Qt* used what some considered to be a non-Free license, and this caused some concern in the Free Software community. However, several years ago, *Qt* was re-licensed jointly under the GPL (GNU General

Public License) and the QPL (*Qt* Public License), thereby removing such objections.

To put the size of KDE into perspective, KDE has around 200 active contributors (programmers, documentation writers and artists), and with over 4 million lines of code in the KDE CVS repository there is a large amount of software around. In fact there is significantly more code than in a recent Linux 2.5 kernel, which contained about 3.7 million lines of source code.

Core components

For quite some time, *Qt* (and therefore KDE) has had excellent support for Bi-directional languages such as Arabic. However, until *Qt* 3.2, *Qt* had no support for Indic languages such as Devangari, Bengali and Tamil, as well as many others. With the release of *Qt* 3.2, *Qt* and KDE developer Lars Knoll was able to demonstrate his work on making Indic languages work in *Qt* and subsequently KDE. Impressively, for *Konqueror* to be able to display web pages containing Indic text, only changes in *Qt* were needed, without any need to change any code in *Konqueror* itself. With excellent support for Western, Bi-directional and Indic languages through *Qt*, KDE is now well suited for use in all areas of the world.

KWin has seen significant rewrites in time for KDE 3.2 by developer Lubos Lunak. At the time of writing, these improvements had just been merged in so that they will be present in the final release of KDE 3.2.



KSVG can show previews of SVG icons in *Konqueror*. Here previews of the *Crystal* icon theme are being shown while the rest of the screenshot uses the *Noia* icon theme.

One of the main aims of the rewrite has been *NetWM* compliance, a standard for window managers published by the freedesktop.org website. Essentially, it specifies a way for compatible window managers to react under certain situations. The standard also means that any *NetWM*-compliant window manager can be used with KDE, by replacing *KWin* with another *NetWM* compliant window manager such as *IceWM*.

Another aim of the *KWin* rewrite has been to significantly restructure the code to make it more flexible. For instance, with the old *KWin* design some bugs were simply unfixable, but

machines that they do not want users to be able to change the settings of, and for those building machines to go in public places, hence the 'Kiosk' name. In KDE 3.0 the kiosk support was fairly rudimentary. Now however there is excellent support for restricting *KWin*, *Kicker* as well as many other programs such as *Konqueror*.

Accessibility Features

After having quite a slow start, much work has been done recently by the KDE Accessibility Project on creating a module of software of accessibility-related applications. Currently the module contains three applications,

“Both KDE and GNOME use the same standard to create application menus and provide application shortcuts.”

thanks to the rewrite this has now rectified. The new *KWin* also has a completely different and far more flexible and easy to use API for developers writing *KWin* styles. Unfortunately, this has the effect of meaning that styles need to be changed slightly to make them work with KDE 3.2. However the benefits that the new API will bring should outweigh this annoyance.

Both KDE and GNOME use the same standard to create application menus and provide application shortcuts. This consists of a file with a `.desktop` extension which specifies the name of the program, the icon to use, and so on. KDE has now switched to a 'VFolder' implementation of the `.desktop` file system – previously the `.desktop` files had to be in a hierarchical structure of directories with each directory having a name which would then become a category on the *KMenu*. Now the category information is held inside the `.desktop` file itself so the `.desktop` file need not be located in any particular location in order for it to be used in the menu.

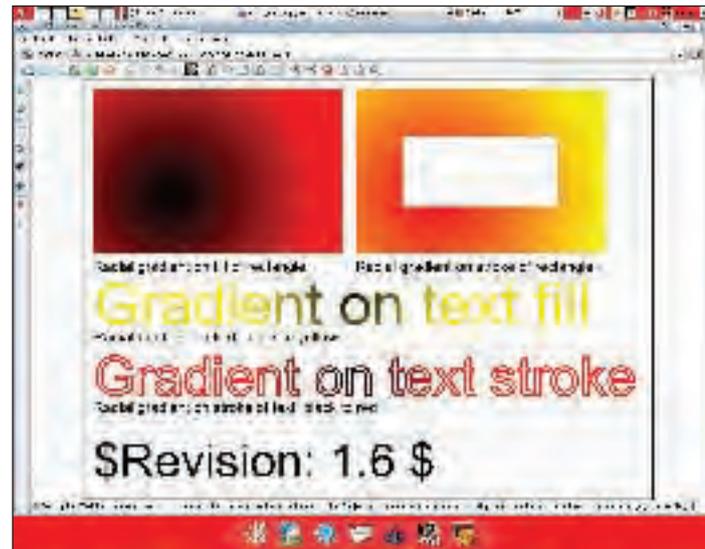
KDE 3.0 first introduced the idea of 'kiosk' support. Essentially this is a KDE-wide method of enabling applications to be restricted in certain ways. This is useful for both system administrators who are building

KMouth, *KMouseTool* and *KMagnifier*.

Olaf Jan Schmidt and Gunnar Schmidt have primarily done the work on this module.

KMagnifier is a tool akin to *xmag* which allows small areas of the screen to be magnified. This will be extremely useful for people that have relatively minor sight problems, since while most of the screen may be legible and visible, it is possible to enlarge just the part of the screen that they are having problems with.

Larger numbers of people than you may think – particularly among the elderly – have trouble controlling the standard computer mouse. Even if a user is able to move the cursor on the screen, they may not have the necessary dexterity in their fingers to be able to click the mouse buttons. *KMouseTool* is a tool which makes it possible to synthesise mouse clicks in the correct position on the screen preventing the frustration that could arise from the inability to click in the right place. >>



Much work has been done on *KSVG* to it can render even the most complicated of images.

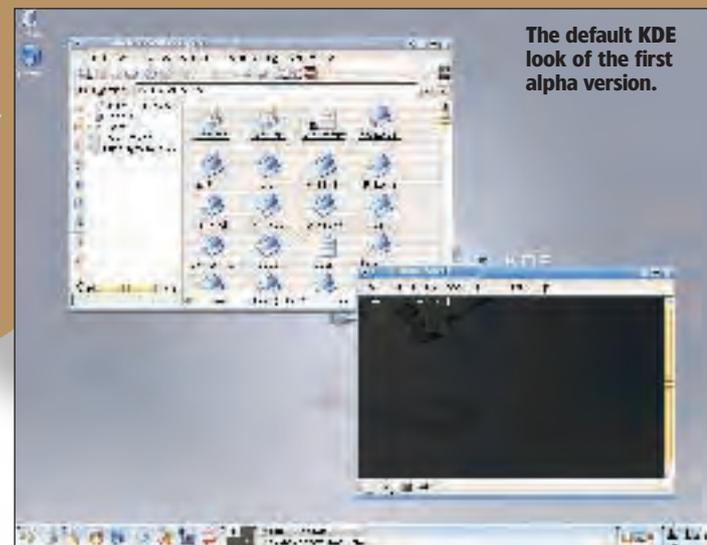
THE BEGINNING

KDE's first steps

The first version of KDE, not surprisingly named KDE 1, was based on *Qt1* and was at the time widely acclaimed for its pioneering work on Free Software user interfaces, winning many awards for technical excellence and innovative features.

The subsequent *Qt 2*-based KDE 2 was also revolutionary. It included a huge number of innovative technologies such as *KParts*, a technology that

allows program components to be embedded into other programs. *dcop*, the *Desktop Communications Protocol*, was developed as a way of facilitating inter-process communication between applications. A large number of applications were also rewritten between KDE 1 and KDE 2. The window manager used in KDE 1, *KWM*, became *KWin*. Additionally *KPanel*, the panel in KDE 1, became *Kicker*.



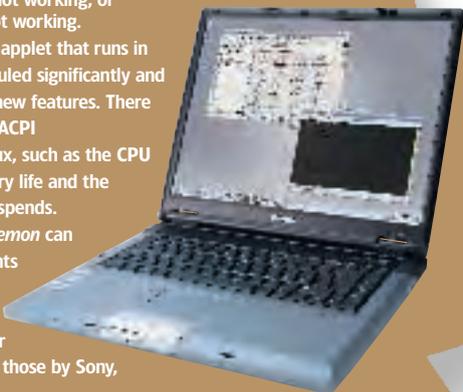
KDE 3.2 PREVIEW

LAPTOP SUPPORT

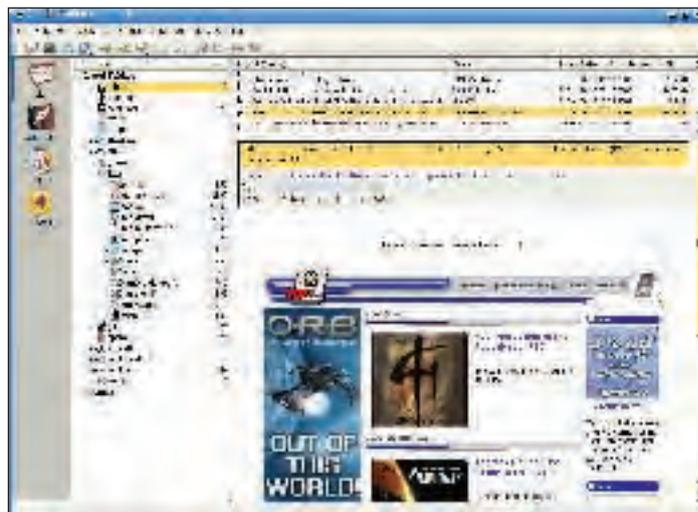
KDE on the move

Sadly, laptops are rather difficult beasts to play with at many times, though the situation has greatly improved over the last few years. While Linux may work well on them, there may be issues such as certain buttons on the keyboard to control the volume not working, or suspend/hibernation not working.

KlaptopDaemon, the applet that runs in *Kicker* has been overhauled significantly and it now supports many new features. There is now support for the ACPI implementation on Linux, such as the CPU throttling to save battery life and the possibility to initiate suspends. Additionally *KlaptopDaemon* can respond to certain events such as the laptop lid being closed. There is also specific support for certain laptops such as those by Sony, Toshiba and IBM.



***Kontact* successfully integrates the calendar and email client into the same program.**



It should also be noted that the KDE Accessibility developers are working with the developers of GTK, Java and *OpenOffice.org* to create a generic accessibility framework known as the AT-SPI (Assistive Technologies Service Provider Interface).

Quanta

Quanta is the extremely impressive web design program in KDE. A huge amount of work has been done on *Quanta* by lead developer Andras Mantia, who is sponsored to work on *Quanta* by Kitty Hooch, the company of one of the original *Quanta* developers Eric Laffoon.

Originally work was being done on a program called *Kafka* to produce a WYSIWYG (What You See Is What You Get) visual web page design tool for KDE. Subsequent efforts to develop this have been merged into *Quanta*, where work is being done to provide WYSIWYG support, based on its *Kafka* progenitor.

Security tools

On Linux, the *GnuPG* program is an implementation of PGP (*Pretty Good Privacy*) and can be used to exchange information which is digitally signed or encrypted. Unfortunately however, many studies over the years have shown that an extremely small number of people understand encryption such as PGP to even a moderately high level. While a program such as *KMail* may allow you to automatically sign and encrypt messages to recipients if you require, public key cryptography requires many other things. Up until now, the main way to do handle these things with *GnuPG* has been via the command line.

Some examples of things that previously had to be done via the command line are the management of people's keys, such as importing and signing. Importing somebody's public key is required before you can send the person a signed or encrypted message. Key signing is the process in which you personally check a person's identity to confirm that they are who they say they are.

KDE 3.2 is the first release to include *KGPG*, a tool for managing keys and a variety of other things

including listing all the people that have signed your key, encrypting directories, and changing various aspects of keys such as trust level.

Personal information management

KDE 3.2 will be an extremely exciting release for people that like programs that have an integrated e-mail client and calendar and so on. Some examples of this kind of program are *Microsoft Outlook* and *Ximian Evolution*.

Several developers have been working on *Kontact*, a program which allows *KMail*, *KOrganizer* (KDE's calendar and scheduling solution) and a few other relevant programs to run from a single main window. Therefore it is easy to switch between the mail and calendar parts of the program and have the calendar interact with the email client in a meaningful way. *Kontact* also has full support for *Kolab*, the software developed by Klarälvdalens Datakonsult and Erfrakon as part of a deal with the German government to develop a free software groupware solution.

Daniel Molkentin, one of the *Kontact* developers, is additionally planning support for making *Kontact* work with *Exchange 2000*, although it remains to be seen whether this will be finished in time for KDE 3.2.

Networking

Kopete, a multi-protocol instant messaging program will be included in KDE 3.2 for the first time. With a large number of enthusiastic developers, *Kopete* supports just about every protocol worth using, including AIM, ICQ, Jabber, Gadu Gadu, Yahoo! Messenger, IRC and plenty more. With a coherent user interface, *Kopete* allows easy communication with your contacts, whatever IM protocol they use. *Kopete* will appeal especially to Linux users who communicate via MSN Messenger – despite Microsoft's recent attempts to dump third-party client compatibility, it seems that *Kopete* works with latest versions of MSN Messenger without a hitch.

KWiflManger, a tool for managing the increasingly popular wireless networking connections, has been added, as has a *Kicker* applet to enable easy monitoring of the strength of the wireless networking signal.

Multimedia

Quite a lot of work has been done on the multimedia side of KDE. Since KDE 2, KDE has used *aRts*, the *Analogue Real Time Synthesiser*, for its sound support. As well as providing a software mixer to allow multiple application to use the sound card simultaneously, *aRts* provides a decoding backend for audio formats such as Ogg Vorbis and MP3. Additionally, for some time it has been possible to use the backend of a popular movie player, *Xine*, to allow *aRts* to decode video as well as audio. This also has the added benefit of automatically allowing *aRts* to support

Primarily this concerns the web browsing part, though much work has been done on file manager part as well with a number of issues fixed.

Konqueror is now able to load *Netscape G/Mozilla*-style sidebars, which have been pretty widely adopted on the web, for example to provide information such as news headlines. For quite some time KDE developers Nikolas Zimmerman and Rob Buis have been working on *KSVG*, an SVG rendering engine. SVG is a technology similar in scope to Macromedia Flash, though has some important differences. Unlike Macromedia Flash, SVG is an official

praise *KHTML* justifying their decision due to *KHTML* being fast and lightweight, amongst other things.

Since then, many of the changes that Apple worked on for *Safari* have been merged back into the main *KHTML* development code so that they are also available to *Konqueror* users. There are still a significant number of things to merge back but these tend to be large things which take significant amounts of developer time and could have unwanted side-effects for *Konqueror* meaning that more time must be spent over it.

It is hoped that eventually both *Safari* and *Konqueror* will be able to share a common *KHTML* code base for their HTML libraries. This however will take significant work on the part of both the original *KHTML* and *Safari* developers. Improvements have been made to the tabbed browsing support in *Konqueror* to make it more flexible and easy to use, and the bookmark

“An extremely small number of people understand encryption such as PGP – in KDE 3.2 the KPGP tool will change this...”

any of the formats supported by *Xine*.

New in KDE 3.2 is *JuK*, an *iTunes*-like program which provides a juke box interface to your music collection. *JuK* has sophisticated play list handling features so music can be ordered into play lists so that tracks can be played in a specific order. Alternatively all tracks can be shown in the same list, or a filtered view based on various search criteria can be used.

JuK is also strongly focussed on the tagging aspects of digital audio, for example ensuring that MP3 and Ogg Vorbis tags have the correct digital tags so that the file name isn't the only thing used for organisation purposes. *JuK* therefore makes it easy to alter the digital tag of a music file. In the future, it will be possible to automatically tag music based on other information already given to the program.

It has not yet been decided exactly what form the multimedia backend will be in KDE 4, but there is a chance that *aRts* will be replaced by something else: technologies like *gststreamer* and X11's *MAS (Media Application Server)* are possibilities that the developers are considering.

Konqueror

A large amount of very exciting work has been going on on *Konqueror*, KDE's file manager and web browser.

standard supervised by the World Wide Web Consortium. As such it is likely that there will be less incompatible changes in the future than for instance may be expected from a proprietary software vendor that is continually trying to persuade its customers to upgrade to the latest version of its software.

In addition, SVG is an XML-based language just like HTML; therefore images can be written simply using a text editor, though for more complex images you will certainly want to use one of the growing number of pieces of graphics software to support SVG.

With the addition of *KSVG*, there will be SVG support in KDE applications such as *Konqueror* and indeed in any other applications that needs to display SVG images thanks to the *KSVG* component.

One piece of news this year which was welcomed throughout the KDE world was an announcement by Apple Computer boss Steve Jobs at the Macworld Expo. This was that Apple's then new browser *Safari* would be based on the HTML rendering engine developed for KDE called *KHTML*. This came as a surprise to many who expected that Apple would develop a browser based on the HTML rendering engine from *Mozilla*, known as *Gecko*. However Apple was quick to

ESSENTIAL KDE COMPONENTS

Qt – The C++ library that KDE is based upon. It provides user interface widgets and more.

Kicker – The program that runs along one edge of the KDE screen to provide access to the taskbar and application menu.

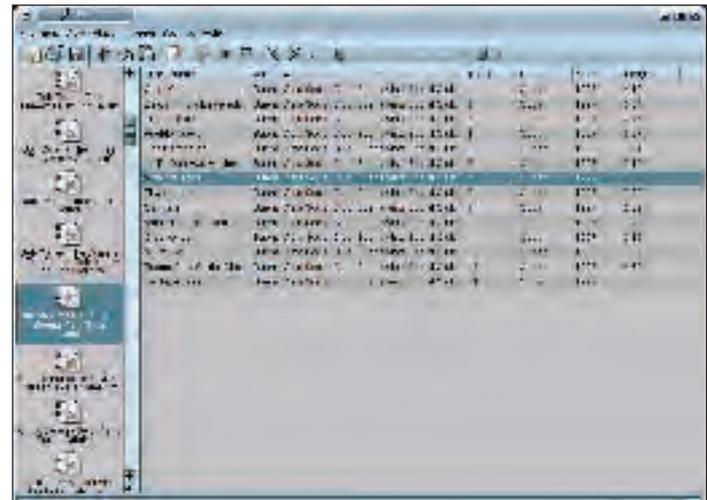
KWin – KDE's window manager, it provides a way of moving windows around the screen and the buttons at the edge of windows for closing and minimise/maximise etc.

KMenu – The application list accessed by clicking on the 'K' button on *Kicker*.

aRts – The Analogue Real Time Synthesiser and audio mixer and decoding back end used to give sound support in KDE.

KHTML – The HTML rendering library developed for KDE in applications such as *Konqueror*.

Konqueror – KDE's integrated web browser and file manager.



handling support has seen large rewrites to improve functionality.

Games & Edutainment

Games are another area of KDE that have not escaped attention in KDE 3.2. The game *KGGoldRunner* has been added. *KGGoldRunner* is a platform-style game where the aim is to collect as much gold as possible.

The fruits of the work by the KDE Edutainment project to create educational and entertainment software originally found its way into the main KDE release in time for KDE 3.0. Since then the Edutainment developers have been hard at work both at adding some new programs,

The JuK audio player has sophisticated play list and filtering abilities comparable to iTunes.



KDE 3.2 PREVIEW

KDE WEBSITES

www.kde.org – The main KDE website.

www.troll.no – The website of TrollTech, the developers of Qt

www.kde-look.org – Excellent site for anyone interested in artwork

<http://accessibility.kde.org> – the home page of the KDE Accessibility Project

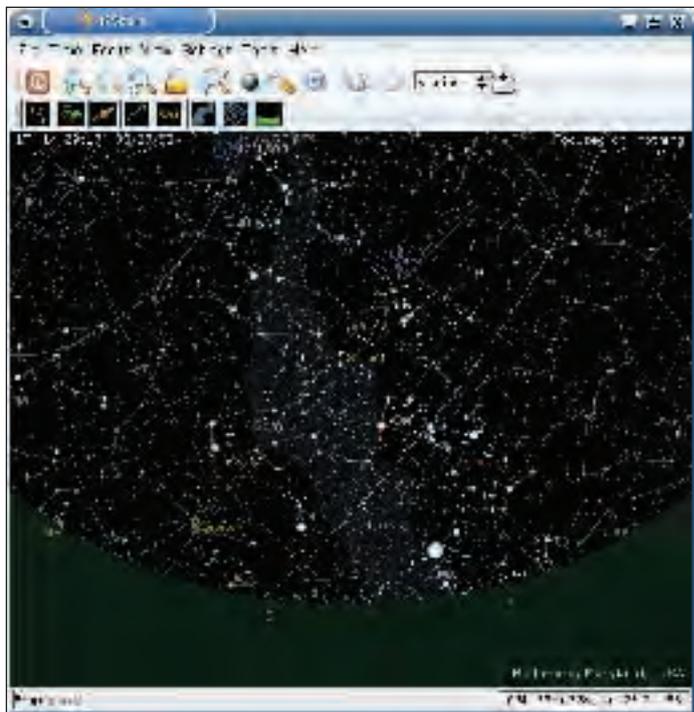
<http://developer.kde.org> – everything for KDE developers and those developing applications based on the KDE libraries.

◀ rewriting limited existing ones and adding many new features.

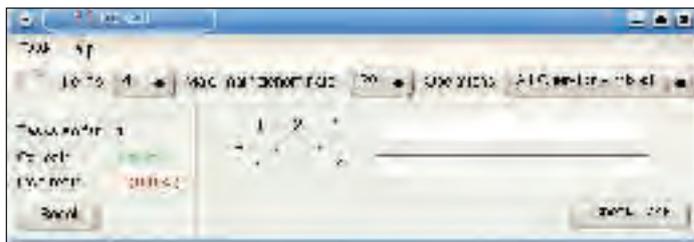
Originally KDE 3.0 had a geometry program called *KGeo* but unfortunately the original developer of *KGeo* did not have any time to work on it and since there were some problems with the program, *KGeo* is being replaced in KDE 3.2 with *Kig*, short for *KDE Interactive Geometry*.

Many improvements have been done on KDE's astronomical software *KStars*, with thousands of new stars being added to the database. Intriguingly, work has also been done to give *KStars* the ability to control telescopes.

Another program to be recently added is *KBruch*. *KBruch* is a piece of mathematical software to help teach fractions. See the web page at <http://edu.kde.org/kbruch/> for more details, and click on one of the listed applications in the left sidebar to get more information on the likes of *KLatin*, *KMathTool* and *KLearnSpelling*.



KStars has a massive database of celestial objects of all types.



KBruch is a tool for teaching fractions.

Art and graphics

KDE 3.1 showcased the *Keramik* widget style, widely acclaimed for its originality. The KDE-Look website (www.kde-look.org) has been used as a gathering place for many artists who work on many areas of KDE artwork including wallpapers, icons, and widget styles. As a result of work first displayed on KDE-Look, KDE 3.2 will introduce a new widget style called *Plastik*. *Keramik* will however stay the default style for KDE 3.2 however and will do so for the foreseeable future.

KDE 3.1 also introduced the *Crystal* icon theme, quite a refreshing new look compared to the old icons which had not changed significantly since KDE 1. Many of the *Crystal* icons have been updated to make them more clear with better colours. In addition, a new icon theme called *Noia* will be imported.

Support for the multimedia features of certain processors has been developed. KDE 3.2 will be able to make use of the Intel MMX and SSE 2 extensions found in Pentium processors. This is used in KDE to speedup certain graphics-related processing and rendering. Support for AMD's 3D Now instructions is also anticipated some time in the future.

KGamma, a tool for adjusting monitor gamma, has been added to enable the gamma of the monitor to be easily adjusted from *KControl*. *KPDF*, a PDF document reader has been imported. Although *KGhostView*, KDE's PostScript viewer, is capable of reading PDF files, the decision to import *KPDF* was taken because *KGhostView* does not support advanced PDF features such as hyperlinks and searching through the text in PDF files.

Support for the XFree86 extension *RandR* has been added to enable the desktop to change resolutions on the fly without the need to restart X11. Special support is needed in Qt and KDE for *RandR* since it is necessary for programs to cope with the fact that the desktop has resized and adjust the size of their windows accordingly.

KWallet

KWallet is a technology to securely store important information such as credit card numbers and passwords

for applications. *KWallet* is in the process of being integrated with *Konqueror* and *KHTML*. Originally touted for inclusion in KDE 3.0, unfortunately a lack of developer time meant that it wasn't finished in time. Since then however KDE developer George Staikos, also responsible for the SSL encryption support in *Konqueror*, has been spending significant time on finishing off the backend and working on the integration with certain areas of KDE.

Therefore by the time KDE 3.2 final is released, *Konqueror* will have the ability to remember form details and save them securely to the hard disk. Additionally integration with *KMail* and *Kopete* is pending and is anticipated to be ready in time for the final release of KDE 3.2.

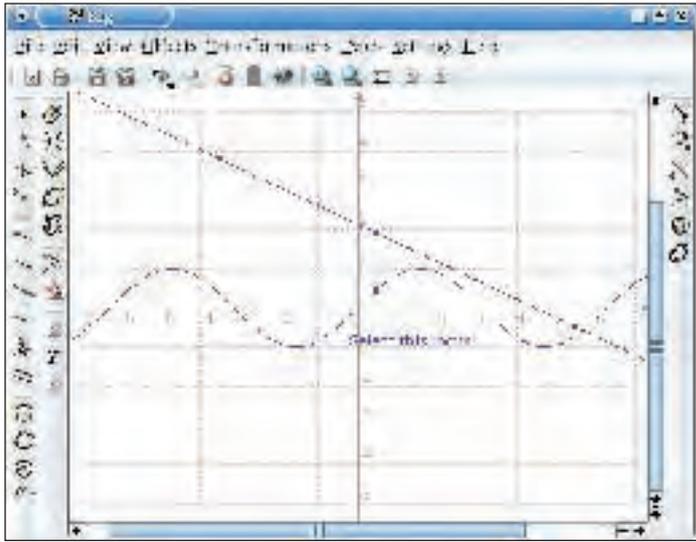
For security purposes *KWallet* has various configuration options such as a Time out. With this feature, a wallet that has been unused for a period of time will ask for the pass phrase again before it will decrypt any more data. Additionally, actions such as starting the screensaver can trigger a request for the pass phrase again.

Development tools

There are two major new tools in the *kdesdk* (*Software Development Kit*) module which are valuable to developers. The first of these is *Umbrello*, a Unified Modelling Language (UML) modeller. UML is a language used in software design. It can do various things such as allow the developer to plan the structure of a program before work is actually started on the code, or to document the code of an existing program.

For example, *Umbrello* can be used with C++ to show inheritance between classes. As well as the ability to draw diagrams from scratch, *Umbrello* can parse existing source code to automatically draw diagrams, making *Umbrello* a good tool for documenting existing code.

KCacheGrind is a more advanced developer tool. It allows software to be profiled and this information visually presented. Essentially, profiling enables the developer to examine how much time the processor spends executing each part of a program so that bottlenecks and other performance issues can be identified.



Kig has many capabilities including Python scripting.

Libraries

With the release of Qt 3, it is widely said that Trolltech more-or-less considers Qt to be complete feature-wise, and that most changes now are concentrated on small API enhancements, bug fixes, and speed increases.

Matthias Ettrich, the founder of the KDE Project and now project manager of Qt development, recently gave a preview of what is in store for Qt 4 at the recent KDE Contributors conference in the Czech Republic.

In Qt 4, the main focus is going to be on separating the currently relatively monolithic Qt into several libraries (rather than having one huge libqt.so file). This should help memory usage since programs will be able to load only the parts of the Qt library that they require. It is possible to write command-line Qt programs which have no need to load the widget libraries for instance. Also considerable work will be done on speed enhancements.

The Qt and KDE libraries work in a very symbiotic manner. In many cases Qt classes are used directly in KDE programs. Other times, thanks to the inheritance and virtual function features of C++, classes are re-implemented to slightly change their behaviour to make them more suitable for use in KDE. Often re-implementing a class takes just a few lines of code.

Qt 3.2 has several new classes. On the user interface side, there is *Q SplashScreen* which makes it easy to implement a splash screen. There is also *QToolBox* a widget that provides a column of tabs.

In KDE 3.2 there are also a considerable number of new classes and libraries. We have mentioned *KSplashScreen*, which is based on *Q SplashScreen* but notably supports KDE's *Xinerama* options. *KAutoConfig*, written by KDE developer Benjamin Meyer, has also been added.

The aim of *KAutoConfig* is to make it very easy to write configuration dialogs for applications. While this sounds like a trivial addition, it is extremely useful. Traditionally, developers had to manually write several lines of code for each setting, to load and save the settings when the dialog was opened and closed.

KAutoConfig removes the need for these lines of code. As well as making it much quicker for developers to write configuration dialogs, *KAutoConfig* is a very valuable addition since it will help to prevent bugs caused by developers needing to tediously write many lines of code by hand.

Of great interest for people from any non-English speaking country is support for International Domain Names (IDN) in KDE. IDN make it possible for any character to be used in a domain name whether it's a French accented letter or from a completely different alphabet.

Since the very start of KDE, development the concept of MDI (Multiple Document Interface) has been frowned upon. It is widely considered to be confusing for certain applications to have multiple documents open in the same window, sharing the same menu bar and tool

bar and so on. However, it has been shown that this approach does actually work for some applications – such as *Konsole*, which can display multiple terminal emulator sessions under the same window. Another application in which it works well is *KDevelop* in which multiple source code files can be opened in the same window.

Qt has fairly poor support for MDI, and as such, the *KDevelop* developers needed a better MDI solution than that provided by Qt. The *KDevelop* authors consequently developed a class called *QExtMdi* which provided a sophisticated MDI solution. Eventually, it transpired that it would be useful for certain other applications to also provide a MDI interface and *QExtMdi* was subsequently moved to *kdelibs* and was renamed as *KMDI*.

Conclusion

KDE 3.2 will soon be released, with the developers working at full speed in anticipation of releasing KDE 3.2 around Christmas-time. After a KDE 3.0 release that largely consisted of switching from Qt 2 to Qt 3, and a

“KDE 3.2 will be able to make use of the Intel MMX and SSE2 extensions that are found in Pentium processors.”

KDE 3.1 release which also had many library improvements though less features, KDE 3.2 will be a very exciting release. There is a vast number of new features, bug fixes and other enhancements to programs for the end user. It also contains a large number of developer-orientated improvements to libraries and a few developer-orientated new programs which will make KDE an extremely interesting platform for both users and developers alike. [LXF](#)

GLOSSARY

SVG – Scalable Vector Graphics, a vector graphics format gaining popularity, particularly on the web.

HTML – Hypertext Mark-up Language, the language that web pages are written in.

UML – Unified Modelling Language, a diagram-based language used by software developed to help develop software.

GPL – the GNU General Public License, the main Free Software used by the Linux kernel and much of KDE.

QPL – the Qt Public License a license developed by TrollTech under which Qt is dually licensed.

PDF – Portable Document Format, a cross-platform method of distributing information which is guaranteed to look exactly the same whatever platform it is being displayed on.

What on Earth is... BITTORRENT?

LXF CD Editor **Neil Bothwick** usually plays his cards very close to his chest, but since it's Christmas, he's sharing this particularly juicy trade secret with us...

» I've seen this BitTorrent thing mentioned on several file download sites, but what is it?

BitTorrent is a file sharing protocol, and the programs to go with it, aimed mainly at distributing large files to a great number of users without placing huge bandwidth and processing demands on a single server.

» So it's a peer-to-peer(P2P) file sharing system? Don't we have enough of those already? What does BitTorrent do that makes it any different from, or better than, the likes of Gnutella, eDonkey, Lopster or Kazaa?

Those other programs are generally used for sharing unspecified

collections of files, everything in your shared directory is available to anyone else using the same system. BitTorrent is aimed more at the distribution of specific, usually large, files, such as ISO images. Other protocols allow you to download a complete file from a single location, BitTorrent will download a single file from a number of other computers, while uploading other parts of it at the same time.

» How can it upload a file before it has finished downloading it?

BitTorrent splits files into pieces, the default size for these pieces is 256KB. When you begin a download of a file, it looks for other computers that have pieces of that file and starts downloading them. As each piece is downloaded, the program lets the network know, so that others without that piece may request a download.

» Surely, if you are all downloading pieces of the same file, you will all have the same pieces?

Instead of downloading the pieces in sequence, BitTorrent looks around to see which pieces have the least copies on offer, the rarest pieces, and attempts to download those. The aim is to ensure that all pieces maintain a roughly equal availability, rather than having hundreds of copies of the first piece on offer and none of the last.

» OK, it sounds interesting. Where do I get it and how do I install it?

You can get it from the cover CD/DVD, in the Magazine/WhatOnEarth directory. The newest version is always available from the BitTorrent website. There are RPM and Debian packages, to be installed in the usual way, and a tarball for those that prefer to install by hand. BitTorrent is a collection of Python scripts, so you need Python installed, and *wxPython* for the GUI program. Installing from the tarball is basically a matter of unpacking it and copying the *.py files to somewhere in your path, `/usr/local/bin` would be a good choice for a system-wide installation.

» Right, I've installed it. Which program do I run to get a list of available files?

You don't. BitTorrent doesn't work like that.

» Then how does it work?

Someone distributing a file via BitTorrent makes it available via a tracker and puts a .torrent file on a web page somewhere. You click that torrent file in your web browser to start the download. Alternatively, you can save the torrent file to your hard disk and pass it to one of the BitTorrent scripts.

» What's this 'tracker'? I thought you said earlier that BitTorrent was a P2P system.

A tracker is a server that keeps tabs on who has what parts of a file available for download. The torrent file contains details of the tracker used to handle a file or collection of files, BitTorrent then asks the tracker where it can get the file from. The tracker doesn't hold any part of the actual file you want to download, it just tells you where to get it. In



this respect, it performs a similar function to the central server of earlier P2P systems, like Napster.

»» So where do I find a list of downloads?

There are many sites that carry lists of downloads and the torrent files needed to start them. Some Linux distros, such as Mandrake and Slackware, make ISO images available via BT. For more a general selection, there is a list of sites containing links at www.dessent.net/btfaq/#links. If you want to try downloading distro ISOs via BitTorrent, there are links to a number of them at www.rpmfind.net/BitTorrent.

»» I clicked on a link, but all it did was download a small file, with a .torrent extension. Where's my download?

There are two ways to handle this. One is to download and save the metadata files (for that is what the .torrent files really are) and then pass them to one of the BitTorrent programs. This isn't as convenient as clicking on a link and starting the download directly, but it does mean you can resume a partial download without revisiting the website. It also means you can browse for files on your dialup connection at home, then take the metadata files to work or university to make use of their bandwidth (please ask the bill payers' permission before doing this). To start the download, type one of the following three commands in a terminal.

```
btdownloadgui.py somefile.torrent
```

```
btdownloadcurses.py somefile.torrent
```

```
btdownloadheadless.py somefile.torrent
```

»» Why the three different commands? Why not just one?

They all do the same thing as far as downloading the file is concerned. The only difference is the way in which they display their progress. The first uses a *wxPython* GTK window to show you what is going on. The second displays the progress of downloads and uploads as text in the terminal. The third command shows no output at all, although it can write the information to a log file. This one is intended for headless servers, and may be the most appropriate one to use when seeding your own files for distribution (more on this later).

»» How do I get my web browser to download the files when I click on a .torrent link?

You need to tell your browser how to handle a file with a MIME type of `application/x-bittorrent`. How you do this depends on the browser. You may not have to do it at all if you installed from a package as the post-install script may well have taken care of this for you. If you do need to do it yourself, adding the following line to `/etc/mailcap` if all that is needed for Mozilla:

```
application/x-bittorrent; btdownloadgui.py %s
```

Konqueror users need to select 'File Associations' from the settings window, click 'Add', set the group to application and type `bittorrent` in the Type name box. Add `*.torrent` to the filename patterns list (this isn't strictly necessary as the web server should send the correct type, but it means you can also launch BitTorrent by clicking on locally saved torrent files) and add `btdownloadgui.py %F` to the application preference list.

»» You mentioned resuming, how does BitTorrent handle that?

Transparently, as it handles most things. BitTorrent downloads in many pieces from different sources anyway, so it makes little difference whether those pieces are downloaded in a single session or at various times. To resume a download, just start it as before, and make sure you give the same destination directory for the download. BitTorrent will see that you already have part of a file and pick up where it left off. There will be some fairly intensive disk activity to start with, as BitTorrent scans the files to see what it's got and what it still needs. Don't worry, this is perfectly normal.

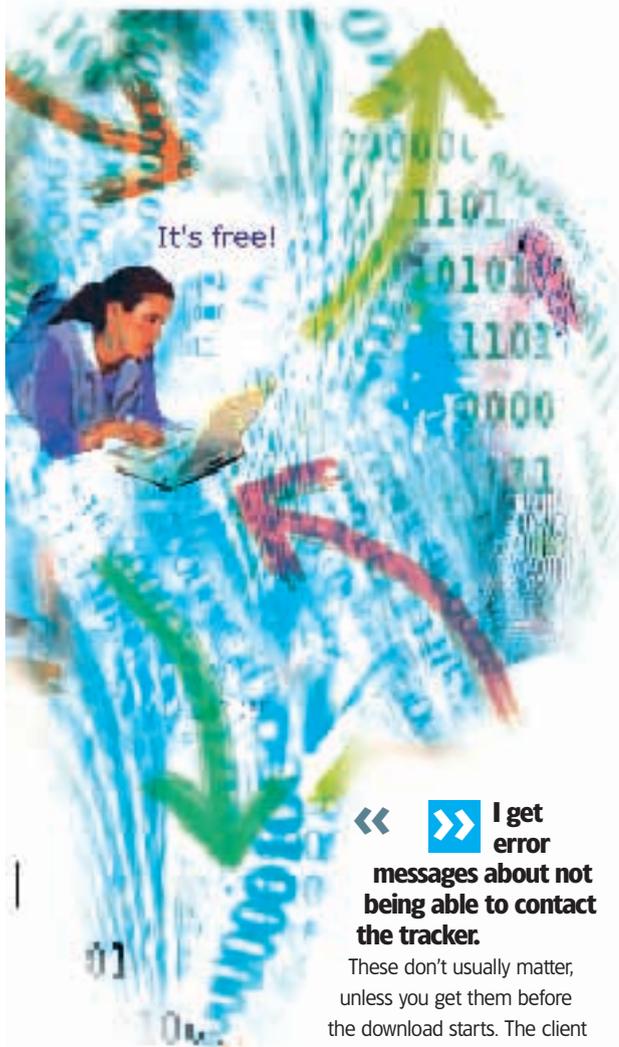
»» I have only just started downloading and already the window is showing upload traffic, what gives?

This is one of the great advantages of the way BitTorrent works. As soon as you have downloaded the first 256KB piece of the file, BitTorrent will make it available to other downloaders. Equally, the other parts of the file will be made available to you as others download them. This means that as the download progresses, and as more people begin to download the file, there are more locations and bandwidth for you to download from. This is the opposite to the way traditional client-server and other P2P systems work, where increased demand usually results in slower downloads for each person. It works particularly effectively with large file, like ISO images, as the number of sources increases within a minute of downloads being started, rather than waiting for anyone to download a full 650MB before sharing anything. »»



Downloading is shared between us

WHAT ON EARTH BitTorrent



« » I get error messages about not being able to contact the tracker.

These don't usually matter, unless you get them before the download starts. The client connects to the tracker periodically to see who else has the pieces it needs and to let it know which pieces it now has.

Once you have connected to another client that has pieces you want, these will download without any need to connect to the tracker, but you won't find out about newly arrived clients or pieces until you can reconnect to the tracker. Although the amount of data transferred between each client and a tracker is tiny, some trackers handle many hundreds of connections and can sometimes become difficult to connect to. This error usually corrects itself automatically over a short time, as the client retries until it connects, meanwhile still downloading from known hosts.

» » If this file has been split up into hundreds of pieces and reassembled for several sources, how do I know it has not been corrupted?

The torrent metadata file contains hashes for each of the pieces as well as the complete file. A piece is not added to the file until it has been checked for integrity, and the whole file is tested when the download is complete. There is no need to include MD5 or other checksum files as the checking built into the process is at least as rigorous.

» » That explains how it works, and how to set it up, but why would I want to download via BitTorrent if other methods are available?

The main reason is the availability of bandwidth. With the normal client-server method, a fixed amount of bandwidth is shared between however many people are downloading from that server. When demand goes up, you either get slower downloads or may find your connection attempt refused. Because every downloader is also an upload source with BitTorrent, as more people download a file, the spare capacity to serve that file to others goes up rather than down. BitTorrent is at its best when demand is high, because more downloaders means more uploaders.

» » How do I stop BitTorrent using all my upload bandwidth?

You can limit the upload rate by adding

```
--max_upload_rate n
```

to the command line, where *n* is the limit in KB/s. You should set this if you are using an asymmetric connection, such as ADSL or cable. Otherwise you could find that your download slows because the upload is saturating your uplink and preventing the download side from sending TCP acknowledgements. Around 80 per cent of your upload bandwidth is considered a good figure. Any lower and your downloads may suffer because you aren't giving enough upload speed.

» » How can I download without uploading?

Why would you want to? BitTorrent is about sharing, so leeching from others without making your own contribution to the community effort is not the way things should be done. BitTorrent enforces this in a way, because it uses a tit-for-tat system of selecting download and upload hosts. The more bandwidth you make available for upload, the better your downloads will be.

» » Can I use BitTorrent from behind a firewall or NAT router?

You can download as normal with no special action beyond the normal installation, although your download speed will suffer for the reasons given in the previous answer. To enable others to connect to you for files you need to configure your firewall to pass through the ports used by BT. Older versions used the range of ports from 6881 to 6889, 3.3 goes all the way up to 6999. The program tries to open the first, then tries successively higher ports until it succeeds. Unless you plan on running a huge number of copies of the BitTorrent client at any one time, the old range of 6881 to 6889 should be more than adequate. How you do this depends on the firewall or router you are using, see its documentation for precise details.

» » How do I continue to make a file available to others after I have finished downloading it?

The BitTorrent programs continue to share a file until you quit them, they don't stop when your download is complete. If you want to keep sharing a file after having downloaded it and quit the program, simply restart the download to the same location. BitTorrent will see that you already have the complete file and make no attempt to download it again, but it will make it available for upload.

There are no hard and fast rules on this, but it is considered good practice to leave the client running until you have uploaded as much as you have downloaded, more if you can spare the bandwidth. If you have an always-on Internet connection, leaving it running overnight is a good way to provide bandwidth to people in other time zones, or those who download overnight, without affecting your own usage of your Internet connection.

» » Why would I want to distribute files this way?

Mainly because it takes most of the load of distributing your files away from you and delegates it to those who have already downloaded (or started downloading it). If all you want to distribute is a small source code tarball and a couple of RPMs, put them on your ISP web space or create a project at Sourceforge. If you have created a CD-sized file, such as a customised version of Knoppix that you'd like to share with others, you'll need rather more bandwidth. As most people only use a minute portion of their available upload bandwidth, why not get them to help you?

» » How do I make a file, or set of files, available via BT?

First you need a tracker. While you can run your own, and may want to for testing, it needs a good reliable Internet connection with a static IP address for public use. Look at the list of links mentioned above, some of these contain free tracker services. Don't just pick a tracker and use it, that's incredibly rude. While many of the trackers are provided for free, you should never use one without permission.

Then you need to create the .torrent metadata file. This file can refer to a single file or a set of files. If you are making a collection of files available in a single torrent, put them in a directory and give that as an argument to `btmakemetatfile.py`. To build the metadata file give one of these two commands, the first for a file the second for a directory

```
btmakemetatfile.py http://mytracker.com:6969/announce myfile
```

```
btmakemetatfile.py http://mytracker.com:6969/announce mydir/
```

There is also a GUI program to do this, called `btcompletedirgui.py`, just run it and fill in the boxes; replacing `mytracker.com` with the name of the tracker you are using.

The torrent file is just a small data file that can be put on a web server and included in an HTML link. The web server needs to be set up to send the correct MIME type with the file. If you are running your own Apache server, add this line to `httpd.conf`

```
AddType application/x-bittorrent .torrent
```

If it is not your server, you may be able to add it to a `.htaccess` file in the directory containing the torrent file, or one above it. If you are using a commercial server, it may already be set up for BT. The easy way to check is to upload the file and click on it. If BitTorrent starts, it works. If you only get an option to save the file, the server is not set up correctly.

Don't forget to let people know about your torrent. They can't download it if they do not know it is there. Once you have the torrent file uploaded, you need to seed the download.

»» You mentioned seeding before, what is it?

The torrent file only contains metadata information on the file, hashes of the pieces and the URL of the tracker. The tracker only has details of who has which pieces available for download. Neither contains any of the file itself. This is a peer-to-peer system, it relies on someone else having pieces of the file for others to download, but right now you have the only copy. So you need to seed the process by making this copy available for download.

Start a download to the current location of the file, as described in the section on resuming a transfer above. BitTorrent recognises that you already have the entire file and immediately gives a "download complete" message, but continues to make the file available for others. Since you presumably want your file to be available all the time, it is probably best to start this download at bootup (or from your dialup script if you don't have a permanent connection) using the headless client. As soon as others begin downloading, they will provide their pieces for others, so the maximum load on your Internet connection will only last as long as it takes for one copy to be distributed. This doesn't mean one person needs to have the whole file, just that each of the pieces is held by at least one of the current downloaders.

»» So you think BitTorrent is the bee's knees, are there any times when it is not suitable?

BitTorrent is best suited to files that have a continual demand, so there are always people downloading or seeding the files to allow others to start their downloading. This makes it ideal for sharing newer, popular files, especially in the period immediately following a new release, when demand can peak very significantly.

BitTorrent is not really suitable for distributing older, archived files, where demand is less popular and likely to be sporadic. A standard web or FTP server is more suitable for files like these.

»» Why doesn't the size of a file grow gradually as it is downloaded?

Because BitTorrent grabs the pieces in an apparently random order, it needs to create the space where the missing pieces will be put. Think of it as BitTorrent first creating a container to hold the file, and then filling it up.

»» If I have a number of files to distribute together, should I put them into an archive or distribute a directory?

Generally, you should distribute files in their usable format. There are two good reasons for not packing the files into an archive. The first is that it makes life that bit harder for those downloading the file as they have to unpack the archive before they can use the files. More importantly, if they unpack the archive to get at the individual files and then delete the archive to save disk space, they cannot make the files available for others to download, as the torrent file only looks for the archive, not its contents.

»» Are there other BitTorrent programs available?

There are other programs in the official BitTorrent package, such as `btlaunchmany.py` that will run simultaneous downloads for all the torrent files in a specified directory. There are also several independent programs, including clients, torrent builders and trackers. Most of these are for Windows, although there is also a Java version that runs on anything with a Java VM. While the official release is considered the safe and stable option, the others often have more features.

»» Where can I find more information?

There isn't much information supplied with the BitTorrent package, although running any of the



clients with no arguments gives a fairly detailed listing of the command-line options. The BitTorrent home page is at <http://bitconjuror.org/BitTorrent/> and there is a helpfully comprehensive FAQ at www.dessent.net/btfaq/. The FAQ is mainly based on the Windows software, but the majority of the information is platform-independent or easily translated. There is also a busy Wiki flourishing at <http://wiki.theory.org/index.php/BitTorrentFAQ>. Yahoo groups, at <http://groups.yahoo.com>, carries several BitTorrent mailing lists; and also using <http://groups.yahoo.com/search?query=bittorrent&submit=Search> will give you a current list of groups that mention BitTorrent, some of which, like *torrent-talk*, *BitTorrentTalk* and *BitTorrent_help* should provide more useful info on all things BitTorrent than we have room for in this brief introduction. [LXF](#)

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 email at linuxformat@futurenet.co.uk or by snail mail, or log on to www.linuxformat.co.uk and post your suggestions in our special forums? Hope to hear from you soon!

Nick Veitch EDITOR

HOW CODE IS REPRESENTED

Including code in magazines can be tricky, but we hope our notation will help it become clear. When lines are too long for our columns, the remaining text appears on the next line in a solid blue box:

```
procedure
TfrmTextEditor.mniWordWrapClick
(Sender: TObject);
otherwise, there is usually a gap between lines:
begin
  mniWordWrap.Checked := false
end;
```

Usually, you'll find the code on our CD/DVD too.

THIS MONTH TEACH YOURSELF...

Beginners' tutorial

Migrating to Linux – remaining backwards-compatible with old files is very important **p62**

HTML/CSS >>

NEW SERIES! Using HyperText Markup Language and Cascading Style Sheets with Linux **p66**

Scribus

Linux Desktop publishing can now be regarded as comparable to using *QuarkXPress* **p70**

Blender

However good your modelling abilities, lots of extra impact can be added to scenes by creating a bit of atmosphere **p74**

The GIMP

Adding depth to 2D images by using shadows is simple, quick and a must-have skill **p78**



Practical PHP

What a lovely PEAR! How the *PHP Extension and Application Repository* solves your PHP problems **p82**

Server School

How to access remotely – for those times when presence of mind equals absence of body... **p86**

TIP OF THE MONTH! NO BELL, NO BELL...

What with the excesses of Christmas, you're probably sick of all the references to bells, whether they're the musical, jingling, door, or tubular variety, so you can be forgiven for being thoroughly sick and tired of dingy-dongy things by the time the festive season is over. So, what's worse than sitting down at your Linux box and having it sound the system bell at you each time you make an error? Here's how to deactivate this annoyance for a peaceful New Year..

There are two systems bells in Linux – one when you're working from the console, and one when you're working inside an X console

such as KDE's *Konsole*. From the console, you can use the **setterm** command to alter settings such as the length and pitch of the system bell, as well as the colours and font types of the text on the screen. To eliminate the system bell from here, the easiest way is to set the length of the bell to 0 by using this command:

```
setterm -blength 0
```

From within KDE, you can stop *Konsole* from beeping in two ways. Under Control Center>Audio>System Bell, you can change the volume, pitch, and duration of the beep – to get rid of the beep completely, just drag the sliders to the bottom. The

alternative (and better) way to disable the beep is to select Settings >Bell from the menubar of *Konsole*, and select 'None.'

This is better than using the *Control Center* primarily because there's another option there, which is 'Visible Bell'. Select this option if you're using a computing environment without any sound-making capability, or work in an office where noises annoy colleagues, but still want an obvious notification if there's a problem, as it will briefly invert the colours in the *Konsole* window whenever it would otherwise have sounded the bell. Bah, humbug! **LXF**

ADOPT A NEW OPERATING SYSTEM

Beginners' Guide to Linux: Migration

Andy Channelle asks: what good is a computer if it can't access the documents and files you've accumulated over the years, and refuses to integrate into a mixed platform network?

Research into the subject has suggested that most novice Linux users are not necessarily new to computers, but are usually migrating from other operating systems. These people bring with them documents, images and practices that must be accommodated.

A shrinking minority of users may experience difficulty finding a suitable application for a specific task under Linux, and for those occasions emulation may provide the solution. There are a couple of ways this can be achieved: either emulate the entire Windows operating system; or find a way of running Windows applications directly under Linux. There are good and bad points about both approaches.

Examples of the former are *VMWare*, (see *LXF43*) which can be installed in either Windows or Linux, virtualises an entire system and will run any x86 operating system; and *Win4Lin*, which is a more limited solution – in terms of host and guest operating systems – and, as such, comes in at a significantly lower price than *VMWare*. The advantage of this tack is that you can install virtually all Windows applications including the latest *Microsoft Office* suite. The disadvantages are that *Win4Lin* currently supports pre-Windows XP systems and will struggle to play most recent games. You will also need a valid Windows 95/98/Me licence and installation disk to use them, and a fairly substantial machine to get the most from it. *VMWare* will install all flavours of Windows but unfortunately its price puts it out of the reach of most home users.

The second option is exemplified by *Wine*, which is an ongoing project to implement the entire Windows Application Programming Interface (API) natively on Linux (see *LXF42* for an in-depth examination). The main project is still regarded as alpha quality, meaning it is not yet considered ready for use on mission-critical applications, but nonetheless is able to run a wide range of Windows software. Moreover, a number of commercial companies have taken the raw *Wine* code and added their own extensions to improve compatibility with specific applications. The most popular commercial distribution of *Wine* is CodeWeavers' *CrossOver Office* which is designed chiefly to run *MS Office* and *Adobe Photoshop*. There are also versions aimed at gamers.

But while emulation can be useful, it's much more efficient to migrate the entire software stack and make the most of the great Open Source applications in the wild.

Application migration

Love it or hate it, *Microsoft Office* has gained a position of ubiquity in businesses and universities across the world with an estimated 400 million users. The chances are, therefore, that before too long, a friend or colleague will hand you a disk or send you an email containing a *MS Office* format file. These are .doc for word processor documents, .xls for *Excel* files and .ppt for *PowerPoint* presentations. A quiet word on the benefits of using less restrictive formats (and this is mostly in the case of .doc files) won't hurt, but remember that most users probably just save things without a second thought to how the data will be stored, so tread softly.

The lineup of Linux applications is growing steadily, but there have been solutions available for a good while that cover the broad spectrum of home and general office use.

MICROSOFT OFFICE *OpenOffice.org* is a free software project that grew out of Sun's donation of portions of the *StarOffice* suite to the Open Source community. It has a core of four applications: word processor, spreadsheet, presentations and drawing; and is considered 'best of breed' in all areas by many reviewers and users. You can download the application from www.openoffice.org.

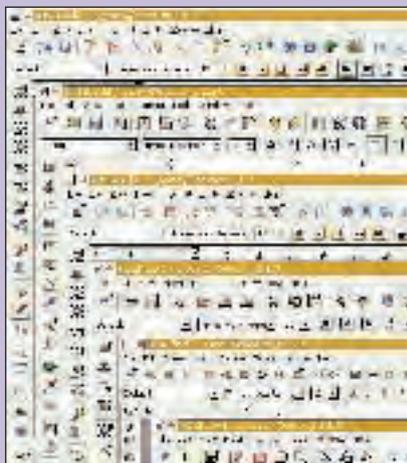
Sun's own commercial offering, *StarOffice 7*, comes complete with SoftwareAG's *Adabas* relational database system, some commercial fonts, a better (much better!) thesaurus and a large collection of gallery items (sounds, web animations etc) and comprehensive app support. It is available directly from Sun at US\$79.00. If you are keen to stick with *OpenOffice.org*, but still feel the need for a decent database, SoftwareAG offers a limited version of its product for download from www.softwareag.com/adabas/ in return for some personal information. The limits don't appear to be too onerous and you can opt out of receiving email from them.

OPENOFFICE.ORG REFURBISHED

Dovetail with your desktop design

One of the small criticisms made against *OpenOffice.org* is that, in order to fit in with so many different desktop schemes, the visuals – icons etc – are very generic. But what if you wanted the suite to fit into your desktop as snugly as *MS Office* fits Windows? Well, help is at hand thanks to the efforts of KDE hacker Rohit Kaul, who has written a tcl script which allows *OpenOffice.org* to fire up with icons sympathetic to your installed theme.

Download the script from www.kde-look.org (it is housed under the themes heading) and extract it. Installation then simply involves opening a terminal, navigating to the correct directory (cd /path) and then typing `tclsh quickInstall.tcl`. Next time you start *OpenOffice.org*, it will be rebuilt with new icons. If you restart and nothing has changed, you probably have the *OOo* Quickstart applet loaded. Unload it and try again.



OpenOffice.org can now be made over to integrate with KDE.

The latest version, *OOo1.1*, will open in excess of 90 per cent of *MS Office* files, and is blessed with some killer features including one-click PDF publishing and support for Macromedia's Flash file format, so you can output your presentations for distribution on paper or on the Internet. File compatibility between *MS Office* and *OpenOffice.org* is not flawless, but then migrating documents from *MS Office 97* to *MS Office XP* is arguably more problematic. *OpenOffice Writer* is also a pretty good WYSIWYG website designer.

The most useful aspect of *OpenOffice.org*, for migrators at least, is that versions of the suite are available for Linux, Windows, Solaris and Apple Mac, so you could install it on both partitions of a dual-boot system and edit your documents regardless of which OS you happen to be using.

The principle alternative to *OpenOffice.org* is *KOffice*, the suite that forms a major part of the KDE project. We'll be taking a closer look at *KOffice* in a future part of the *LXF Beginners'* series.

JASC PAINTSHOP PRO/ ADOBE PHOTOSHOP

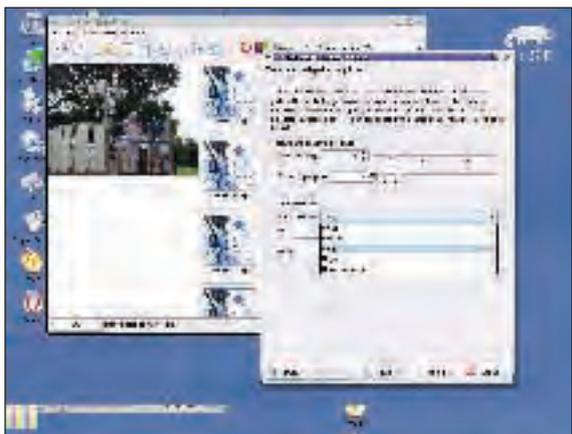
Digital images will pose less of a problem thanks to the forward-thinkers who opted to publish format specifications rather than hide them away. Common image formats are JPEG (.jpg), Portable Network Graphic (.png), TIFF (.tif) and GIF (.gif) and you should have no problem accessing these. If you have used *Adobe Photoshop*, *Corel Photopaint*, *PaintShop Pro* or – shudder – *Microsoft Paint*, you will be pleased to note that one of the unqualified successes of Open Source development is *The GIMP* (*GNU Image Manipulation Program*). This is an image editing and creation package capable of everything from playing with family photos to building quite complex website interfaces.

The current version of *The GIMP*, available in virtually every distribution, is 1.2. The next major release is rumoured to be imminent and adds the more familiar window-based menu system that users of other programs have been requesting.

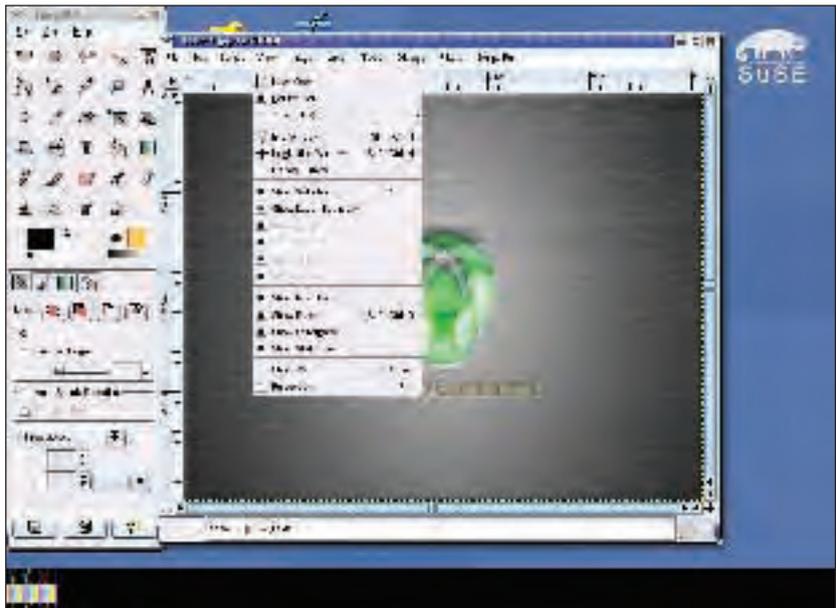
APPLE IPHOTO

Reflecting the trend to put the computer in the 'centre of your digital universe', Mosfet, creator of the wonderful High Performance Liquid theme engine for KDE, has developed a photo management and display application called *PixiePlus*.

PixiePlus is included with many distributions – it can also be downloaded from the sporadically available www.mosfet.org website (if you can't find it here, search www.rpmfind.net) – and is similar in scope to *ACDSee* or *Apple iPhoto*. As well as



***PixiePlus* can help you manage your family snapshots and publish them on the Internet.**



being able to display some 80 different graphic formats, *PixiePlus* is a capable slideshow application and is great for taking screenshots, but its killer feature is the ability to take a directory full of images and create a gallery web page which can be posted straight from the application onto the Internet. Smart!

The latest version of *The GIMP* has a more traditional user interface option.

WINAMP/iTUNES

If digital photo-management can be thought of as a nice extra, organising your music collection is rapidly becoming the big thing for computer users, and as hard-disk sizes expand, ripping every CD you own can quickly become an all-consuming passion. There are a number of applications to extract music from CDs and encode them as MP3s or Ogg files including *Grip* and *KaudioCreator*. The standard MP3 playback application is *XMMS*, though if you really do have 300 or so CDs already ripped, *JuK* (www.slackorama.net/cgi-bin/content.pl?juk) which is planned for inclusion in KDE 3.2 has some useful tools for managing vast collections. The next version of *JuK* will use the *MusicBrainz* system (www.musicbrainz.org/) to provide information on CDs and MP3s that you're listening to.

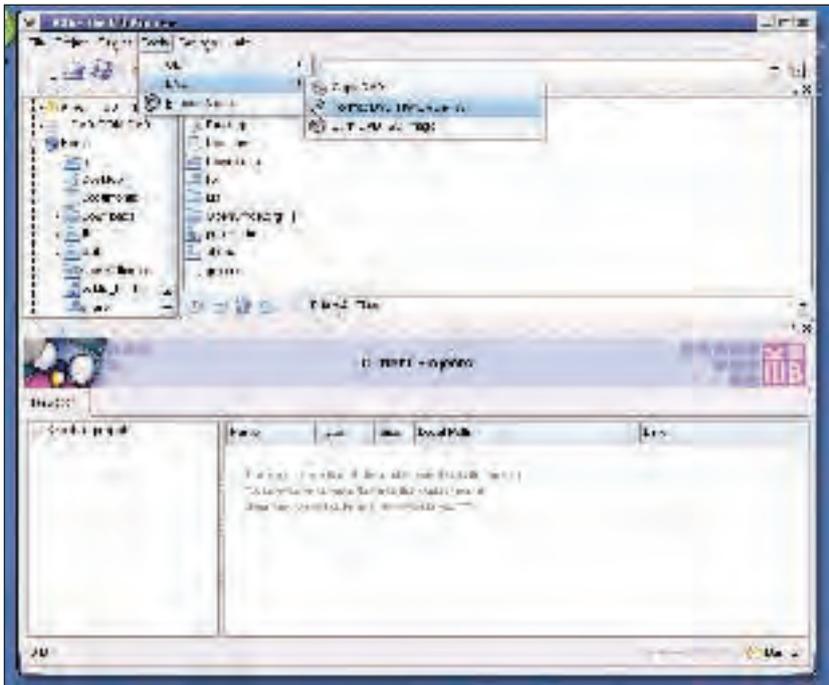
NERO BURNING ROM

Lots of PCs come with CD burners, and Linux has a great compilation application called *K3b*. It is capable of creating audio, data and mixed CDs and, with a little training, can also handle



***XMMS*, or is it *iTunes*? This is a very skinnable application.**

TUTORIAL Beginners' Linux



K3b's latest roster of features include tools for burning DVDs.

VideoCDs and DVD. *K3b's* strength lies in its modular nature, plug ins exist for most tasks and almost everything relies on drag and drop (including dropping things from *Konqueror* straight into the 'files' window). The application integrates incredibly well with KDE so, for instance, right-clicking on a directory offers the option to 'Create a data CD with K3b..'

WINDOWS MEDIA PLAYER/WINDVD

Putting aside the issues of legality surrounding the playing of commercial DVDs on Linux, there are applications available which will do the job. Those that are included in mainstream distributions, however, are 'crippled' to avoid breaking the law. The two main video players for Linux are *Mplayer* (<http://mp.dev.hu/homepage>) and *Xine* (<http://xinehq.de>), both which can be coaxed into playing *Mary Poppins* or *The Matrix Reloaded*. See their respective homepages for more details. Usefully, *Mplayer* is capable of using Windows codecs so, for instance, with a correct installation, you can access the many Quicktime videos on the Internet (film trailers etc.) without having to boot into Windows or buy an Apple Macintosh. Again, we will cover this in more depth in a future tutorial.

INTERNET EXPLORER

Linux is blessed with the three best browsers available: *Mozilla*/*Firebird*, *Opera* and *Konqueror*. Any one of them is capable of handling your browsing needs more than adequately.

OUTLOOK

There are many alternatives to *MS Outlook* that will improve your security, privacy and efficiency. Before settling down, test *KMail*, *Ximian Evolution*, *Thunderbird* – the *Mozilla* offshoot – and *Opera's* integrated *M2* client. All have their good points, but *Thunderbird* has by far the best spam filters, which will make a big difference to your mailing habits, especially if you want to protect your children from those ubiquitous, unsavoury offers.

SAMBA

We've covered setting up native networking recently, but Linux is equally at home working within a mixed architecture network. In

SAMBA RESOURCES

More information

Samba is a complex subject, but there is plenty of information out there if you need to do something beyond the very basic. The first place to start is www.samba.org which has a small selection of well targeted tutorials on subjects ranging from sharing printers to setting up complex permissions.

Information and advice on using *Webmin* is available from www.webmin.com, and you can find more about *SWAT* (*Samba Web Administration Tool*) on the main Samba site.

If you want to get deeper into the system, there is a very good tutorial at www.samba.netfirms.com/index.htm which starts nice and easy and then adds a little more complexity.

Pick up *LinNeighbourhood* from www.bnro.de/~schmidjo/

fact, recent benchmarks have shown that the Open Source software designed to link Linux and Windows computers, *Samba*, is more efficient as a file and print server than the most advanced Windows server OS.

Just like the Network File System (NFS) that we've already covered, *Samba* is a client and server system. If you are using a Windows XP laptop to access files on a Linux desktop, you'll need to install the server on your Linux box. If, conversely, you intend to access files on the laptop from the desktop, you'll need to install the client. Best to take the kitchen sink approach and install both! If you are using Red Hat or earlier versions of Mandrake or SUSE (by which we mean previous to the current versions, 9.2 and 9 respectively) it would also be useful to install *LinNeighbourhood*, a small application that reproduces the experience on using Network Neighbourhood on Windows.

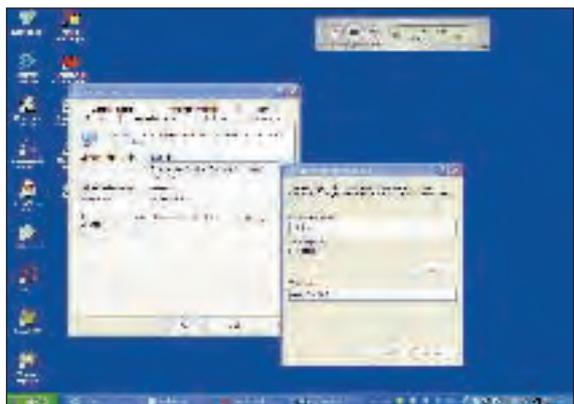
As before, for the purpose of this tutorial, we are going to assume that your network hardware is set up and working and you can ping freely between machines.

Sharing

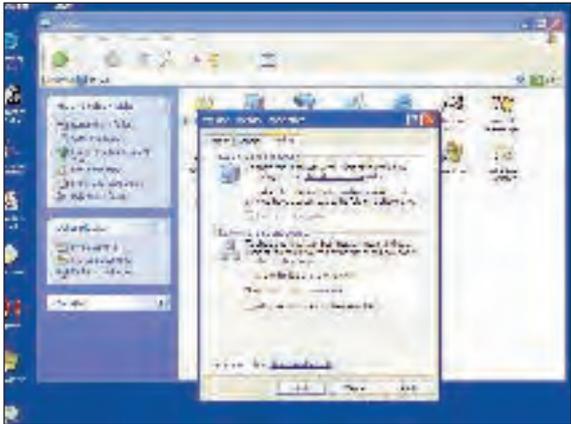
To ping from a Windows XP machine, do Start>Run and type **cmd** in the dialog. This will open what used to be called a 'DOS box' on the desktop. Using ping here is almost the same as using it in Linux, though the letter used to set the number of echo requests is different.

```
ping 192.168.0.1 -n 10
```

This will send ten packets across the network and tell you if any go missing. Computers in a Windows network belong to a Workgroup, by default it will be named WORKGROUP – the name is case-sensitive. To find and change the name, open the



Microsoft's SMB uses a scheme of machine names and workgroups to identify elements within the network.



This is where you define *Samba* share in WindowsXP.

start menu, right-click on 'My Computer' and select 'Properties'. You'll see information revealing the name of the computer and the workgroup it belongs to. You can change either by hitting the 'Change' button.

Next up is defining the files and folders you actually want to share. To do this open *Windows Explorer*, navigate to the folder you want to share, for instance My Documents, and right-click on it. Select Properties and choose the sharing tab. To allow access over the network, highlight the 'Share this folder on the network' radio button. In the text box, you can give the share (which is what we will call the *Samba* accessible folders from now on) a name different from the folder name if, for example, you need it to be more identifiable. You will also find a crude form of Linux's permissions in this dialog, giving the option to 'Allow network users to change files'. Leaving this blank will allow read only access to the share.

Once you hit the 'OK' button, the selected folder becomes an SMB share and will be available to users across the network.

In Linux setting up a share is a similar process, once you have the *Samba* server installed. Go into your usual package management tool and search for the server. Install it. Configuration of *Samba* is quite a black art and there are many ways to set things up, using graphical clients or by writing directly into `/etc/samba/smb.conf`.

As the most universal method for configuration, we'll have a crack at doing it manually, which means becoming the root user or, in KDE, typing

```
kdesu kwrite
```

in the Run dialog, and opening the file in this fine text editor. Scary bit first: `smb.conf` has over 100 user definable parameters, fortunately the file should already be populated with a pretty usable set of defaults. In fact, there's a good chance that the only thing you will need to do is change the workgroup name to, in the case of our small network, WORKGROUP. The process of defining shares is done through *Konqueror*.

In the current versions of Mandrake (on the LXF47 CD/DVD last month) and SUSE there are dedicated tools to configure *Samba* within the Control Panel/*YaST* utilities. You can also install *Webmin* which has a useful configuration tool. See the Resources box on the previous page for more details.

If it's not configured to start at bootup, you may have to manually launch the *Samba* server. Open a console and type:

```
insserv smb
```

You can set *Samba* to launch at run time by editing the Run Level in either Mandrake Control Panel or *YaST*.

In KDE you should now be able to open *Konqueror*, right-click on a file, select Properties and get access to the 'Local Net Sharing' dialog. The first job here is to hit the 'Local File Sharing' button. Access to this dialog needs the root password so enter it and select 'Allow users to share files from their HOME directory'. Hit 'OK'. You'll surmise from this that users can only configure sharing for files within their own home directory.

Next, right-click on a directory again. You should notice a new entry in the menu: Share. The Share dialog (really just an extra tab in the Properties box) has three well signalled options: Not shared, Shared – read only for others, and Shared – writable for others. Chose the second or third option to make the directory available across the network!

And that is the configuration done – straightforward stuff! The shares from both the Windows and Linux PCs should be available through any *Samba* client.

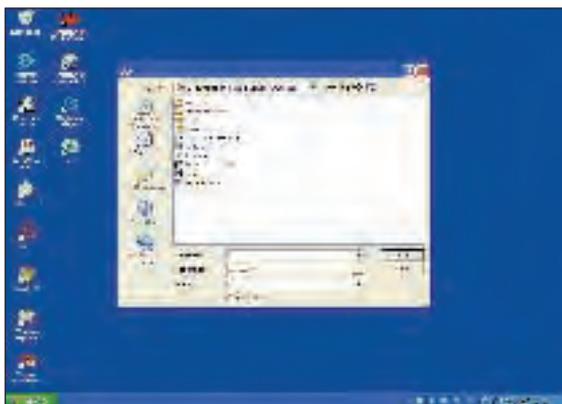
Accessing

With both machines networked, you should now be able to go into My Network Places in Windows XP and select 'Map a Network Drive'. Now hit the 'Browse' button and navigate to Microsoft Windows Network/Workgroup/Linux/<share name>. The 'Linux' portion of the URL will be whatever 'name' your Linux box has been assigned. Hit 'OK', assign the share a drive letter so you can access it from within any Windows application, highlight the 'reconnect at logon' radio button and click on 'Finish'. You may be asked for a user name and password the first time you set this up. If this is the case, be sure to use your full user name, for instance 'Andy Channelle' instead of 'andy' and use your Linux account password. Connecting to your *Samba* shares from Linux should be even simpler, but there are occasional problems. So, we'll start with the easy way.

Open *Konqueror* and in the location bar type `smb://` to access the WORKGROUP folder. This directory provides access to *Samba* shares on every machine in the network, so on our two PC network it contains subdirectories called Linux and Notebook, with shares available within.

Nautilus users can similarly type `smb://` to gain access to shares.

If *Konqueror* and *Nautilus* aren't playing, you may have more luck with *LinNeighborhood*, which is a self-contained *Samba* client which works in a similar fashion to Microsoft's Network Neighborhood. With *LinNeighborhood*, you can mount *Samba* shares as a specific user (it's just the same as mounting a digital camera or other removable storage device) and then access them in the normal fashion. [LXF](#)



Not a blue screen of death – *StarOffice* for Windows can access *OpenOffice.org* for Linux documents via *Samba*.

NEXT MONTH

Christmas is a time for giving, and what better way is there to show your love for your offspring than to take digital photos of them now in order to have some material with which to embarrass them at some point in the future? So, next issue we'll be looking at digital images – getting them in, mashing them up, getting them back out again. We'll also try our best to squeeze in a short tutorial on configuring *Samba* as a print server so you can print via Linux across your network.

WEB PROGRAMMING

Using HTML/CSS with Linux

PART 1 Jono Bacon kicks off a new three part series on getting to grips with HyperText Markup Language and Cascading Style Sheets.



Few people would disagree that the rise of Internet has been one of the most significant technological advancements of our time. The Internet has risen from a small network empowered by secrecy and people with wispy beards, to a huge global network where anyone and everyone can connect, bearded or not. Although the Internet is largely thought of as the World Wide Web (WWW), the Internet actually consists of many different networking technologies.

The birth of the WWW and its rise as the most popular of Internet technologies has largely been down to its creators' (the W3C) determination to keep all aspects of the WWW open and accessible to anyone. This has resulted in a number of open standards being developed that dictate how to communicate on the

Internet and how to display content. The different open technologies all focus on particular aspects of accessing information on networked computers, and some of these technologies include:

TCP/IP Low level communication protocol for communicating between machines

HTTP A protocol that specifies how to access hypertext documents

HTML A language for creating web pages

CSS A standard for styling web pages

XML A standard for creating markup languages such as HTML

XSL A standard for styling XML documents

RSS A standard for creating information feeds for websites

These are only a few of the many open standards and languages that have been developed under the wing of the W3C, and we will be looking at two of them; HTML and CSS.

The structure of web pages

HTML (HyperText Markup Language) is a special language that has been developed to allow anyone to create web pages. This special language includes a number of special features that allow you to specify fonts, create tables, add text, show images and many other elements that are used to create a web page. Although considered a language, HTML differs from typical programming language in the sense that it is a tag-based language. HTML contains a number of tags that indicate how something should be formatted in the web page. An example is:

```
<i>This text is italic</i>
```

In this example we have two tags. The first tag is `<i>` and this indicates that we will to make some text italic. All tags are formed from some text inside angled brackets. The letter *i* inside angled brackets is an Italic tag in the HTML specification (the set of rules that dictates which tags do what). The text following `<i>` in our example above will be changed to italic, and the `</i>` tag will end the area to be made italic. The `</i>` tag is referred to as a Closing Tag and is formed from the original tag, but with a backslash / before the tag contents. In the above example, anything between `<i>` and `</i>` will be made italic.

Now we are familiar with a typical tag, we will look at a slightly more complex example:

```
<p align="left">This is my paragraph.</p>
```

The `<p>` tag is a paragraph tag in HTML, and is used to separate paragraphs in a web page with some whitespace. Each paragraph in the page would use the `<p>` tags to be displayed

BROWSER DIFFERENCES

Broken standards

One of the problems with the WWW is the fact that although the W3C has set out how HTML and CSS should behave, not all browser manufacturers stick to those guidelines. Many of these companies have added additional features or not implemented standard technologies such as HTML and CSS fully.

Although we could suggest that you ignore these non-compliant browsers, the most popular web browser currently used on the Internet (*Microsoft Internet Explorer*) is one of these browsers that is not 100% in implementing the W3C guidelines properly. The problem with some browsers (*Internet Explorer* included), is that they don't demand that the programmer writes proper HTML and CSS, and still work if the programmer makes a mistake such as leaving off a closing tag. Although great in the short-term, this makes all compliant browsers that enforce the programmer to write proper HTML break when they encounter this sloppy code. It is therefore recommended that you make your web pages using only compliant HTML and CSS and write it properly.

HTML RESOURCES

Learning HTML, like any other language, is helped with a good set of resources and tools. Some of the following websites and tools will be useful:

www.w3c.org

W3C Website

This is the center-point for the development of Internet technologies including HTML. There is a huge amount of documentation on this site.

<http://validator.w3.org>

HTML Validator

Use this website to check that your HTML is valid and approved by the W3C.

www.w3schools.com

W3Schools

This is an invaluable website with a number of free tutorials and resources for learning Web technologies.

<http://quanta.sourceforge.net>

Quanta web editor

KDE-based web development editor.

<http://bluefish.openoffice.nl>

Bluefish web editor

GTK-based web development editor.

correctly. The difference with this example is that we have an additional part inside the `<p>` tag that specifies `align="left"`. This additional text is called an Attribute and is used to tweak the tag more specifically. These attributes are not available for all tags (such as the `<i>` tag) as not all tags need further refinement. The align attribute in this example is used to set the justification of the text in the paragraph. If we set this attribute to left as we have done above, we will get left aligned text, and other settings can include center, right and justify.

Getting Started

With the simple introduction to tags out of the way, let's get started and actually make a web page. To do this you will need the simple tools of the trade that include:

- A web browser (*Mozilla, Konqueror, Galeon* etc)
- A text editor (*Kate, gEdit, Vi, Emacs*, etc)

These are the basic two tools that you will need to make web pages using HTML and CSS and virtually all the other technologies described earlier. The other tool that you will possibly need is a web server if you wish to put your pages on the

Internet, but for learning HTML and CSS, a text editor and browser is all that is needed.

Your particular choice of editor is totally up to you, and you may want to use one of the specific web development editors such as *Quanta* and *Bluefish*. As for a web browser, you should be aware that browsers sometimes show HTML a little differently (see the *Browser Differences* box below for more on this), but I would overall highly recommend *Mozilla*, and specifically *Mozilla Firebird*, as the browser of choice. Mozilla is specifically built around the open standards and is a very high quality browser.

To get us started, create a file called `firstpage.html` and put the following code in it:

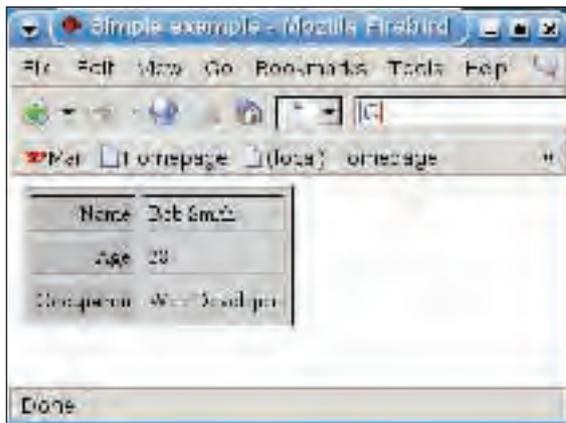
```
<!DOCTYPE HTML PUBLIC "-//W3C//DTD HTML 4.01//EN"
"http://www.w3.org/TR/html4/strict.dtd">
<html>
<head>
<title>A first webpage</title>
</head>
<body>
<p>This is the content of my webpage!</p>
</body>
</html>
```

This code contains the main elements of a simple web page. The first tag is known as the Document Type Definition (DTD) and specifies the type of code that is underneath it. In this example we are using the HTML 4.01 strict DTD that ensures we write HTML that is strictly parsed (in other words, we are going to write proper HTML). More information on DTDs is in the Validating HTML box over on page 68.

The next line in the file contains the `<html>` tag. This tag, and the closing tag at the end of the file indicates that all content and tags between the two tags is HTML. It is common to have the

TUTORIAL HTML

Fig1 Colour codes may look daunting at first, but you can always use *The GIMP* to generate the correct code for the exact shade that you want to use.



◀◀ **<html>** tag at the start of the file and the **</html>** tag at the end, as the file generally contains entirely HTML. The first tag inside the **<html>** block is a **<head>** tag. The head block contains tags that set information about the file. In this block we have a single tag which is the **<title>** tag. This tag specifies the title name of the document and usually appears in the top window bar of the browser.

After the **<head>** block is the **<body>** block that represents the main page in the browser window. Inside the **<body>** block you can see that we have a paragraph block.

Developing our page

So far our little webpage has only a single, solitary paragraph; Amazon.com and eBay can rest easy in their beds tonight! To add some content, we need to make use of some more HTML tags. The first thing we may wish to do is use larger and smaller font sizes. These can be used with the six different levels of headings, with **1** as the largest and **6** as the smallest:

```
<h1>huge</h1>
<p>normal text</p>
<h6>huge</h6>
<p>normal text</p>
```

Adding pictures to our web page is simple with the use of the **** tag. This tag looks like this:

```

```

This tag specifies an image name using the **src** attribute. The **alt** attribute is used to give a textual description of the image, and is particularly important for blind web users who use screen readers. There are other attributes such as height/width for setting the size of the image and border for setting the thickness of the border (**0** is no border, and then a number greater than **1** is the thickness of the border).

VALIDATING HTML

Writing 'proper' HTML is commonly known as writing 'valid' HTML. The validity of the HTML is determined by the HTML specification written by the W3C. The specification is available in slightly different formats, and these are specified in the first line of the HTML file in the DTD. The DTD describes what kind of code the file contains so it can be validated.

Validating HTML and other W3C technologies can be done by going to <http://validator.w3.org/>. You can get the W3C validator to check your code and let you know of any problems with it. When there are no problems, your code is certified Valid HTML and should work in any reasonably compliant browser. It is always advisable to write valid HTML.

Adding Tables

Tables are a common method of showing organised information using lines and table cells to make the information easier to read. HTML includes some special tags and attributes that can make creating tables simple. Here is a simple example of a two-column table with three rows:

```
<table>
<tr>
<td>Name</td>
<td>Bob Smith</td>
</tr>
<tr>
<td>Age</td>
<td>28</td>
</tr>
<tr>
<td>Job</td>
<td>Web Developer</td>
</tr>
</table>
```

In this example we have the outer **<table>** block that indicates that inside the block we will be building a table. The next tag is a **<tr>** tag that indicates that we are creating a row in the table. Inside each **<tr>** block we have two **<td>** blocks. The **<td>** tag lets you create a table cell. The table is built as the code is read in, so the first **<tr>** and **<td>** definition refers to the first row on the table.

There are some additional attributes for the **<table>** tag that allow you to fine-tune its appearance. This includes border for setting the outlines of the table. If you set border to **0**, there will be no border, and if you use a number you will create a border that varies in pixel thickness. Not only does **<table>** have attributes, but **<tr>** and **<td>** have their own set of attributes. An example table that uses these attributes is:

```
<table cellspacing="5" border="1" cellpadding="3"
bgcolor="#d3d3d3">
<tr>
<td align="right" bgcolor="#bebebe">Name</td>
<td>Bob Smith</td>
</tr>
<tr>
<td align="right" bgcolor="#bebebe">Age</td>
<td>28</td>
</tr>
<tr>
<td align="right" bgcolor="#bebebe">Occupation</td>
<td>Web Developer</td>
</tr>
</table>
```

You can see how this example looks in Mozilla in **Fig1** above. The cellspacing, border, align and cellpadding attributes may seem fairly simple to understand, but the **bgcolor** attribute may look a little strange to you; after all, what does **#d3d3d3** and **#bebebe** mean? These are called 'colour codes' and are hex references to colours.

These codes are split into three sections; red: **#RR_____**, green: **#__GG__** and blue: **#_____BB**. Each colour can range from 0 to F (0, 1, 2, 3, 4, 5, 6, 7, 8, 9, A, B, C, D, E, F) and refers to the colours intensity. As an example, **#FF0000** will produce a completely red colour as the red part is as high a value as possible, and the other two parts (green and blue) are at their

OPEN SOURCE WEB

Ever since the adoption of the WWW, developers have had easy access to the source code of websites. If you look in the View menu of Mozilla for example, you will see a View Source option that will show you the HTML of the web page. This feature (which is available in other web browsers also) is incredibly useful for learning HTML as you can see pages that you like and see how they work. Use this feature wisely and it can really help your understanding of HTML and its concepts.

lowest values (00). Using the three parts we can combine them to create combinations of the three colours to produce virtually any colour you can imagine. For an equal mix of all the colours, try #888888. To bypass this rather arcane method of specifying colours, you can use the colour picker in *The GIMP* to select a colour and it will show you the HTML code.

Connecting pages together

One of the most glorious aspects of the WWW is that geography becomes no issue. You can visit a site in Norway and click on a link and be taken to a site in Australia instantly. These links allow us to weave connections between our site and other sites, and it is using these links in which search engines harvest their databases. Creating a link uses the Anchor tag:

```
Go and visit the <a href="http://www.linuxformat.co.uk">Linux
Format </a> website.
```

In this example we are specifying that the words *Linux Format* in the sentence are a link to another website.

We do this by wrapping the words inside the `<a>` tag that has a href attribute that specifies the name of the site that we are linking to. Not only will the link tag allow the text to be clicked on to go to the *Linux Format* website, but the link text will be changed to a different colour. Links can also be applied to images.

When creating links, the format of the link address is important. If you are linking to a normal website, you must include the http:// part so the browser knows that you are referring to a website as opposed to another link. You can link to another file in the same directory by specifying just the name of the file, for example:

```
<a href="mypage.html">link</a>
```

You can also link to an email address by prefixing the email address with mailto; for example:

```
<a href="mailto:my@emailaddress.com">mail me here</a>
```

An example

To put our knowledge of HTML together in a practical example, I have created a very small web page that displays information about *Linux Format* magazine. This page displays the Linux Format logo, some information, a table and a copyright notice. The code is as follows:

```
<!DOCTYPE HTML PUBLIC "-//W3C//DTD HTML 4.01//EN"
"http://www.w3.org/TR/html4/strict.dtd">
<html>
<head>
<title>About Linux format</title>
</head>
<body bgcolor="#fff79e">

```

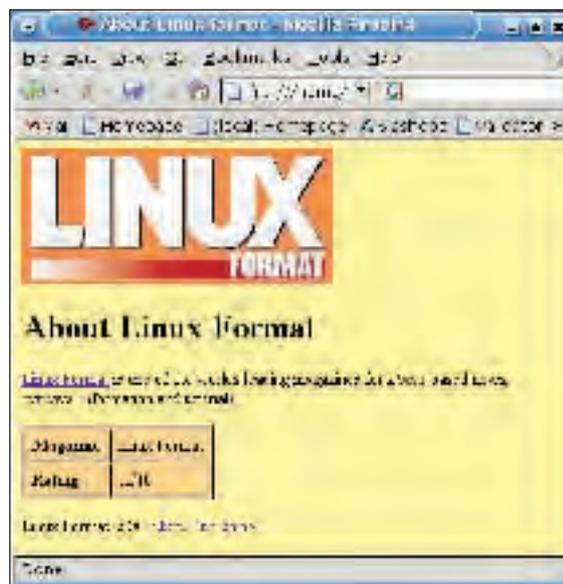
```
<h1>About Linux Format</h1>
<p>
<a href="http://www.linuxformat.co.uk">Linux Format</a> is
one of the worlds leading magazines for <i>Linux</i> based
news, reviews, information and tutorials.
</p>
<table border="1" cellpadding="5" bgcolor="#ffd78d">
<tr>
<td><b>Magazine</b></td>
<td>Linux Format</td>
</tr>
<tr>
<td><b>Rating</b></td>
<td>10/10</td>
</tr>
</table>
<p>
Linux Format is &copy; <a href="http://www.futurenet.com/">
Future Publishing</a>.
</p>
</body>
</html>
```

The first thing to point out in this example as shown in **Fig2** is that this code is *not* W3C valid code. This code is merely an example to show use of combined HTML tags; we will be fixing its validity with CSS in the next issue.

Summary

In this last code example, we can see how we have used the **bgcolor** attributes to change the colour of various parts of the code. We have also made use of the **** tag to make text bold, and we have used a special HTML code **©** that will display a © copyright symbol. To top this example off, we have two links to the *Linux Format* and *Future Publishing* websites.

In this issue we have taken our first step into the world of HTML. We have looked at some example tags and created some simple pages, but we have not really even scratched the surface of HTML, let alone CSS – the scope of this first tutorial is just to introduce the subject. Over the next couple of issues, we will explore the world of web programming a little further to inspire you to dump your website building apps and get coding! **LXF**



NEXT MONTH

We will be looking at using HTML and CSS to create page designs and layouts that can house our information. Until we resume next issue, it is recommended that you have a look at the source code of websites that you visit (see the *Open Source Web* box) to see the different tags and blocks of code. Do not let these complex websites put you off, but looking at the code will get your eyes used to the layout of HTML for more advanced cases.

Fig2 The end result of the code presented on this page of the magazine.

DESKTOP PUBLISHING

Scribus 1.0 – the basics



Finally, Linux users are blessed with a DTP application! Andy Chappelle conducts a tour of its features.

With our increasingly over-designed culture, two things stand out when it comes to imparting your message: simplicity and consistency. The former breeds clarity and accessibility in your work, while the latter helps to initiate and foster a relationship with your reader. Fortunately, *Scribus* puts a suite of tools at your disposal that makes creating clear and consistent documents less of a challenge.

When you create a new document (File>New) you have a number of decisions to make including the size and orientation of your pages, margins (which will be expressed in the 'Default Unit' value) and the first page number. You can opt to add an automatic text frame with predefined columns and gutter (the space between columns) width, and also determine whether your document will feature facing pages (for creating double page spreads) and whether it will begin with a left or right page (especially useful if you're designing your cover in a separate document and want to begin on page 2 or 3).

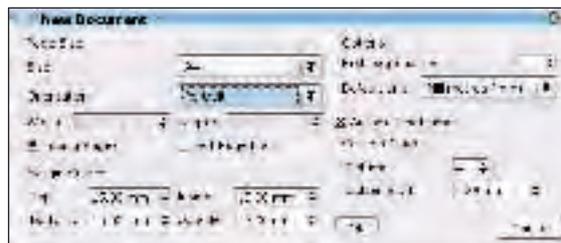
Your new document will open with a single page. To add more, do Page>Insert... Select the number of pages you want to create, where they will go (after page x etc.) and whether they will be based on existing templates. We have opted to stick with the default 15mm border on each edge of the page.

Stamping ground

For the creation of publications such as newsletters, templates are an absolute necessity, allowing you to build your pages on a consistent framework and saving the time and effort involved in adding recurring features and setting up columns, gutters and

HOLD THE FRONT PAGE!

The Scribus Toolbar



Scribus will automatically set up columns for you.

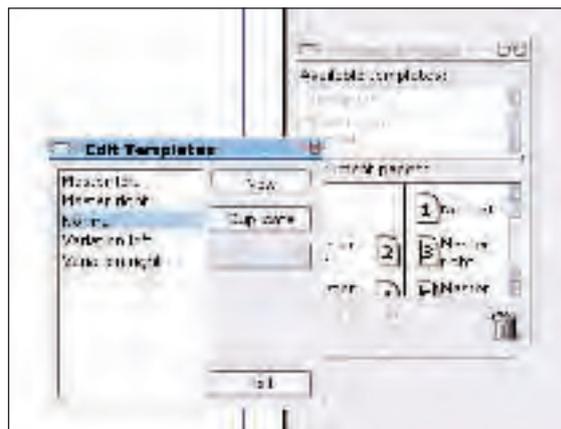
margins for every page. The act of creating the template will also force you to think about your presentation.

The Template dialog is accessed via the Edit menu (initially it will contain only the 'normal' template). When opened, you are dropped into template mode and everything you enter onto the blank page will automatically be added to the selected template.

To create a new template, select 'New' from the dialog. Give your creation a name and a left/right designation and start adding content. For a newsletter, you will need at the very least one template each for your standard left and right pages. You may also want a variation set with a different number of columns or an asymmetric design. To save a little time you can 'Duplicate' your master template, rename it and make the changes necessary for your variation scheme.

You can add any content to a template including guides (simply drag and drop them from either the vertical or horizontal rulers at the edge of the screen), images, repeating text such as headers and footers or background colours. Once you have finished polishing, hit the Exit button and the template will be replaced with the standard document view.

You can apply templates to existing pages by opening the Page Palette (Tools>Show Page Palette) and dragging a template name – ie Master page 1 – from the 'Available templates' palette down onto the right page number below. To quickly edit a template, double-click on its name in the 'Available templates' palette. Now you are ready to begin the fun part: designing your pages... Well, nearly!



Creating good templates will save work in the long run.

Get some style

in this tutorial, we're creating a short run literary fanzine of 16 pages designed to be printed on an inkjet or laser printer and then photocopied for distribution. *Scribus*, of course, is capable of CMYK output suitable for printing at a professional bureau, but that is way beyond our needs at this stage.

Before we begin bringing in content, we will first define text styles to cover most eventualities – headlines, crossheads, bylines, captions, body text *etc* – and also set a few colours (or greys in this case) to function as our standard palette. Both Style and Colour dialogs can be accessed via the 'Edit' menu.

While there are a few text formatting options you can use from within the main application, most of the paragraph options – first line indent and space before or after paragraphs – must be set in the Styles menu.

The Styles management dialog box is fairly self-explanatory with the New, Edit, Duplicate and Delete buttons doing as you'd expect. The really useful aspect of creating and using Styles in your publications is the ability to make global changes to the look of your document without having to change every individual element.

Our first style is body text: Times New Roman, 11points (pt) on 13.2pt leading (the gap between lines), aligned left. We've set the first line indent to 12pt.

The colours dialog allows us to define our corporate colours and, if necessary, tweak them after the pages have been laid out. In a black and white document, it may seem frivolous to set up a specific colour for devices such as boxouts (at 30 per cent black) – we could just add the colour shade individually – but if, on proofing, it becomes apparent that the boxouts are too light for white text, we would only have to change one definition, instead of going through the whole document replacing the parameters for each box individually as we go along. Remember to save your styles and colours once you finished in their respective dialogs, otherwise they will just disappear into the ether.

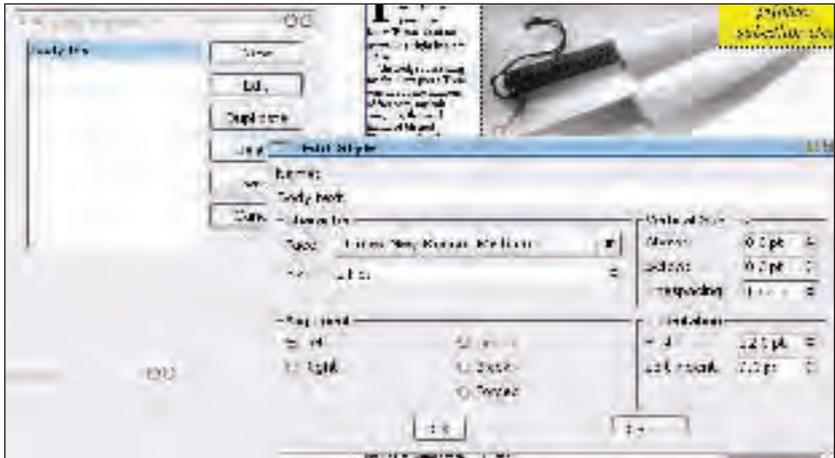
Bring it on!

Desktop Publishing applications such as *Scribus* only have a limited range of design and text editing facilities, the generation of these elements is the province of applications such as *The GIMP*, *OpenOffice.org* or some other specialised programme.

Our front page features a masthead (the title of the fanzine) created in *The GIMP*, a stock image with a little processing, a headline and body text for the main feature and some 'puffs' for other parts of the fanzine.

Scribus has quite a selection of image filters and is capable of handling everything from JPEG and TIF to PNG, MNG and GIF. To bring in our masthead, we create a picture box using the 'Insert Picture' tool (sixth icon from the left on the toolbar – see box, right), and then right-click on the box and select 'Get Picture'. This brings up a standard file selector where we can choose a picture, in this case 'masthead.png'.

Once imported, you may find the image is far too big for the box on the page. To change this you can either right-click the image and select 'Modify' or open the Measurements Palette (Tools>Show Measurements). Check the 'Scale to Framesize' radio button. The 'Proportional' button does as you would expect, but you can reduce or enlarge the image in the box by selecting 'Free scaling' and defining position and x/y scaling. With Free scaling selected, you can also move the image around within its boundaries with the content/hand icon on the toolbar. If you are manually altering size but you want to keep your image in



The Styles dialog gives access to more comprehensive paragraph formatting options.

proportion, remember to click the small icon to the right of the scale values to show an unbroken chain.

Once the masthead is in place, we can start adding text boxes to the page that will contain the main body of the article. Our front page design entails a story title across the centre with three text boxes to contain the first part of the text – the left-most box bisected by the title. In the space above the title we will have an illustration relevant to the the main text.

As with the image, we can right-click a text box and select 'Get text.' However, as the design stands, the text will just flow into the first column. We need to create a chain of text boxes.

With the first column selected, click the Chain icon in the toolbar, and then click anywhere in the second column. Now, with the second column selected, click the Chain icon again and then click on the third column. The fourth column on the far-right will function as an index and so will not be part of the chain.

Next to the Chain icon is the Unchain icon. To use this we would first need to select a pair of linked text boxes, using the standard click, shift-click method, and then click the icon to disconnect them.

Now the columns are linked, importing the text should see it flow across your page.

There is a bug in the current version of *Scribus* which means selecting text across linked textboxes is impossible, so if the text is in need of heavy editing or formatting, it is best to create one large text box on a new page (or even a new document) and do the work in their before cutting and pasting across.

Remember *Scribus*, though *au fait* with lots of graphic format, likes text as plain text. However, you can copy and paste directly from *OpenOffice.org* or even, if you are using KDE, drag a .txt file from *Konqueror* into a text box.

With our text imported we can begin fitting it, configuring paragraphs using our predefined styles (click an insertion point in the text and select the relevant Style in the Measurements Palette) and editing the copy to fit the space.

R-r-r-runaround

A drop cap is the large initial capital letter that commonly appears at the start of a feature. Unfortunately *Scribus* doesn't have an automatic drop cap feature yet, so we'll have to create one manually (see the box overleaf for an alternate method).

First create a small square text box, add your initial letter – in our case a T – in the correct typeface, and alter the size so it covers three lines of your normal text. Position the box so that the top of the capital letter lines up with the top of your first line



TUTORIAL Scribus

◀◀ of text. The problem you'll notice is that you may have one or two blank lines below the drop cap. To get round this, select the text box with the drop cap in, open the Measurements Palette and select the 'Shape' element. Ensure that 'Text flows around box' is not activated.

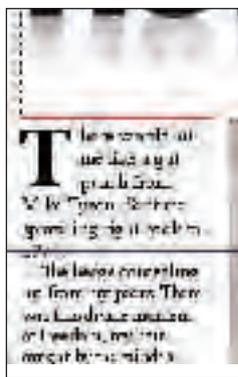
Now the problem is that the text runs underneath the drop cap, so now select the rectangle drawing tool from the toolbar and create a small box over the letter T. In the colours palette, select 'None' for both background and outline. Select the box with the drop cap in and do Item>Bring to Front.

Creating a drop cap in this fashion is a useful introduction to the concept of text runaround, which is one of the essentially elements of DTP. *Scribus* can run text around rectangles, ellipses or irregularly shaped objects with user definable gaps between the object and surrounding text.

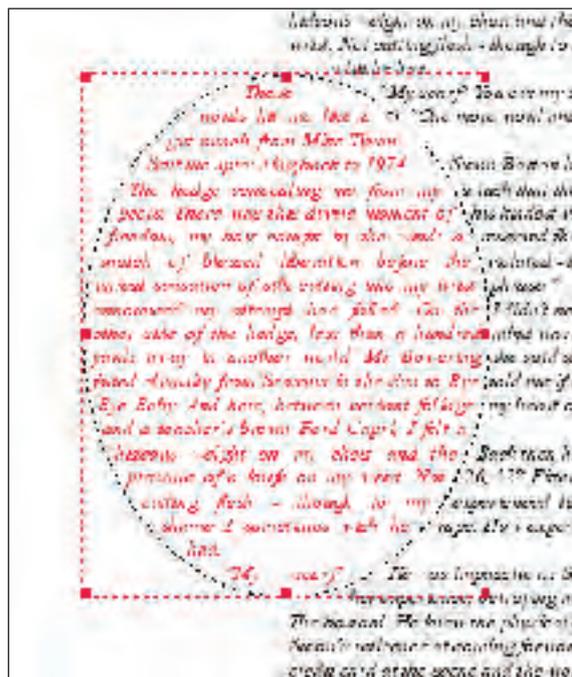
To experiment, open a new document, create a rectangular text box and fill it with a text file. Now create an ellipse shaped text box and fill that with text (if you want to see the effect well, use different coloured text). When you move the ellipse text box and intersect the first one, the text will flow around it. In the Measurements Palette you can define the distance (in points) between the box borders and text.

You are not, of course, limited to just square or round boxes. As well as being able to create polygons from scratch, or adding and editing nodes to existing shapes, you can turn a piece of text – obviously big, fat typefaces work best – into a text or picture box, by selecting its box with the standard selection tool and doing Item>Convert to Polygons. It is then simple to ungroup the text (Item>Ungroup), transform individual letters into picture boxes (right-click letter Convert to...>Picture box), and then import your image.

If this still isn't enough, and you need even more control over a shape, select it, open the Measurements Palette and select the 'Shape' tool. Hitting the 'Edit Shape' button offers node level editing tools similar to those in *CorelDraw*. Any shape is possible. Beware though, the Undo feature doesn't appear to work during this sort of editing: save often just in case you completely mess up.



Drop caps can be fiddly to sort out, but by no means impossible.



Run text around irregular objects with ease.

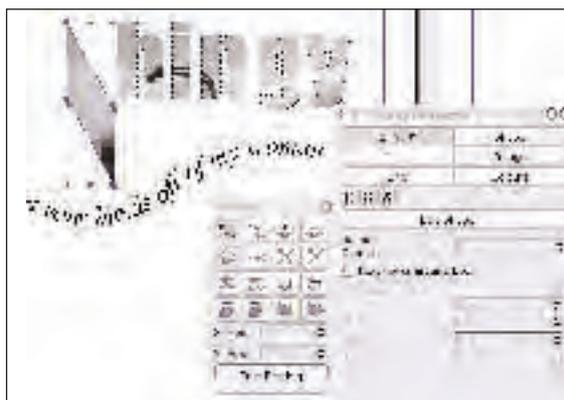
ALTERNATIVE DROP CAP

Managing runaround

A simpler, though less flexible method for creating a drop cap, is to first create the letter in *GIMP*, save it out as a .png file and bring it in to *Scribus* as an image file. The advantage of this is that for letters such as **W** you can contour the runaround if you don't want an unsightly gap at the base line of the capital letter. The disadvantage is that changing the first word of your feature involves more effort!

Scrap it

As well as following a particular template style for laying out text, almost all publications have recurring elements that are used on more than one page. In the case of Linux Format, issue-on-issue as well as page-by-page, there are Top Stuff Award logos, On The



Shape and box designs can indulge your every whim.

LAYERS

A multitude of uses

One of the most powerful features of *Scribus* is its layers, which can be used for a number of purposes. The primary use, as the name suggests, is for creating multi-layered designs, perhaps 'locking down' specific elements of your page furniture such as page numbers or watermarks, while having the freedom to move foreground parts without disturbing the rest. Layers can also be used for multilingual design, with each language housed on its own layer, or to add post-it style notes all over your publications for printers or others who may need access to your work.

If you've used the layers feature in applications such as *The GIMP* or *Adobe Photoshop*, the dialog will be instantly familiar.

A row of icons at the base of the dialog allow you to add, delete or move layers up and down. In the main part of this window is the list of layers. Here you can rename a layer – just double-click its name and type over the old one – and determine whether each layer is visible (the small eye icon) or will print (the printer icon).



GIMP users will feel at home with *Scribus*'s layers palette.



Scrapbook has many uses, and is easily portable too.

CD and On The DVD logos, verdict boxes, etc. To create or import these devices from scratch is a time-consuming process, so the Scrapbook facility in *Scribus* – while seeming a little superfluous at first – quickly becomes an essential utility. Its great strength is the fact that the Scrapbook is not just restricted to containing graphics and you're not restricted to having just one scrapbook.

To add something – text, image, or a group containing both – select it, right-click to launch the context sensitive menu and choose 'Send to Scrapbook'. Once in the Scrapbook, you can do File>Save As... and save your scrapbook in a new .scs file.

Though not designed for this purpose, the Scrapbook can also be used as a movable colour palette. To do this, create a square and fill it with your chosen colour, then send it to the Scrapbook and save it.

When you open the scrapbook again – even on a different machine – you can drag and drop the square into your new document and the custom colour will be added to the palette.

No print proofing

So finally it is time to output our efforts and the format of choice these days is Adobe's .pdf (*Scribus* can also output .eps and .svg files). There are quite a few options to work your way through, some of which work better than others.

To get the dialog, do File>Export>Save page as PDF... By default *Scribus* will dump its PDF output into your working directory. You can change this by typing a new path in the address bar at the top of the PDF dialog and hitting 'Change'.

Some things to look out for:

COMPRESSION We experienced real display problems when outputting files with compression enabled. It's best avoided until a later release has cleared up any lingering problems.

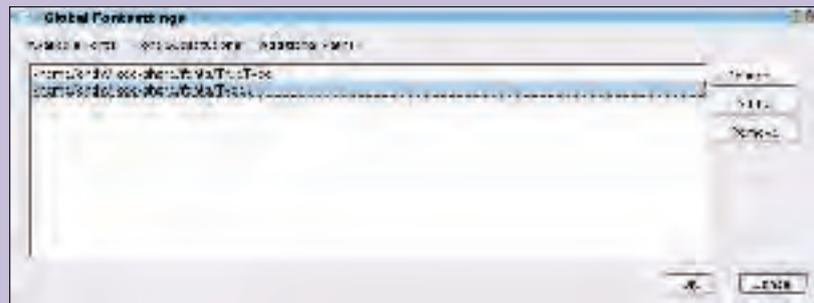
DOWNSAMPLE IMAGES Useful in creating screen proofs for emailing. Downsampling to 72dpi shouldn't affect screen rendering.

COMPATIBILITY *Acrobat 5.0* provided the best results. If you expect to archive your work on the Internet, selecting 'Save linked text frames as PDF-articles' makes good sense.

FONTS embedding fonts increases file size, but for the sake of an easy life, it's the best option. If you're happy to leave the rendering of the more common faces to chance, remove them here.

IT'S FONTS YOU WANT

Tweaking the typography



If you know where all your fonts are located, adding them is easy.

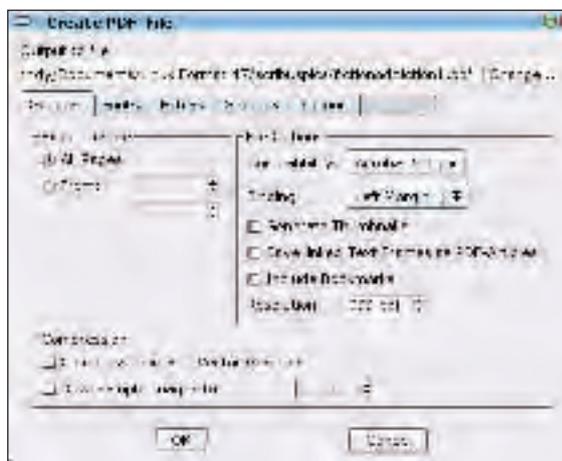
Good quality fonts are one of the keystones of good page layout and, by coincidence, one of the areas where Linux has traditionally lagged behind its proprietary cousins. Fortunately, the font catch up is well underway. Linux – and hence *Scribus* – can use any TrueType or Type1 fonts. Moreover, type foundry Bitstream has made a series of fonts available for Linux users under an Open Source license, which means no more relying on the pretty-but-proprietary Microsoft TrueType fonts.

You can pick up the fonts at www.gnome.org/fonts. Simply download and extract the archive to your fonts directory and install them as normal.

To make fonts available to *Scribus*, you will need to add the font location to the 'additional paths' dialog.

In order to access this portion of the fonts dialog (in Edit>Preferences>Fonts) before you open a document. Select the 'Additional Paths' tab, hit the 'Add' button and browse to your fonts directory. In our standard SUSE 8.2 installation, KDE's font were housed in /home/<user>/kde/share/fonts/ with a directory each for TrueType and Type1 fonts. If you're dual-booting you can easily add a path to your Windows fonts.

When it comes to design, one of the usual hallmarks of an inexperienced user is the kitchen sink approach to type. Just because you have 1001 TrueType fonts on your PC, doesn't mean you have to use them all! Most documents will look best with just two – or at a push three – typefaces. You can use different weights and devices such as bold or italic to add emphasis.



Some of the PDF options are more successful than others.

EXTRAS this section of the output dialog features tools for creating standalone presentations. Multi-page slideshows can be created complete with page transitions.

SECURITY allows you to password protect files and limit the uses of the final product, such as preventing printing, copying or altering parts of it.

COLOUR Select whether the document is intended for print or screen consumption.

Scribus offers tools just beyond the sophistication of *QuarkXPress 3* for free, so is a potential killer app for Linux as a whole that (with further refinement) could eventually be as important as *The GIMP* and *OpenOffice.org* in establishing Linux on the desktop. [LXF](http://www.linuxformat.co.uk)

ILLUSTRATION BY TOMASZ WOJCIWICZ - www.blender3d.org

FINISHING TOUCHES

Getting shady with Blender



PART 9 This month **Jono Bacon** introduces us to simple techniques to increase the atmosphere of your *Blender* scenes.

Over recent months we have been looking at a number of different parts of the *Blender* system to develop our skills through different challenges that we face within the field of 3D scene composition. This month we are going to take a slightly different approach and add some atmosphere to our scenes. Although we will be learning some new features in *Blender*, we will primarily be looking to consolidate what we know already and use these techniques to create feeling in our scenes. It is best to use this month's article as a means to develop a sense of realism in your imagination which can then be modelled more effectively in *Blender*.

Never fear the atmos

The first question we need to define when attempting to create atmosphere is what atmosphere actually is. The definition of atmosphere tends to vary between artists, and people interpret it differently, but it basically boils down to the immersion of the viewer in the scene. As an example, if we are playing a computer

game such as *Resident Evil* or *Silent Hill*, the atmosphere of our intrepid hero wandering through deserted streets is amplified by use of darkness, light, chilling sound effects and other subtle aspects that would be present if you were really there. Imagining you are inside the scene is the key to creating this atmosphere, and your challenge is to recreate what you would see if you were right there in the thick of the action.

Earlier, I stated that people see atmosphere in different ways, and this is all dependent on your ability to imagine the scene as it should be, and recreate these small elements effectively; precisely how this is done is dependent on how you view the scene. We will be looking at a few examples of how to create some effective looking scenes, so you should pick some hints and tips that can be applied to your scenes also.

Atmosphere foundations

Atmosphere in 3D modelling comes from three main areas – materials, light and modelling. Modelling is generally the first port

of call for developing your feel in the scene. If, for instance, you are creating a gothic-style setting, you would want to include the typical decor of that style of scene. This would probably include lots of spikes, flaming torches, swirling structures and harsh dark statues. The second main aspect that is usually applied then is the materials that form the objects. For this same gothic example, you may use lots of concrete, metal, granite and other harsh and dark materials. The final major area is the lighting. The lighting of the gothic style scene will likely be very subdued and gloomy, with varying areas of shadow and light. These three major areas combined will convince your brain into separating the reality of you as a viewer and draw you further into the scene. Remember that no single area will develop your atmosphere, but all of these areas are combined to varying degrees.

Another issue that you should take into account when developing your scenes is the size of your scene objects/actors and their relation to each other. As an example, if you want to create a scene that is awe-inspiring and huge, create lots of small small actors and a huge set for them to be in. An example of this is in action films where there are thousands of army soldiers heading towards each other. You can really create epic-looking scenes by using size to your advantage. Remember that as a 3D modeller, you control everything and do not have the typical limitations involved in hand-built physical scenes.

The importance of light

Of the three types of atmosphere inducing elements, light is possibly the most critical. You could take possibly the most intricately modelled scene with impressive textures and materials, but if it is lit incorrectly, the scene will not convey atmosphere at all. As an inverse effect, you can take a simple scene and give it incredible atmosphere by effectively lighting it.

We will explore some effects by taking a simple room scene and applying lights. If you create a new *Blender* scene, and remove the default plane, you can then add a cube to the scene. Select the front four vertices and press the **X** key and select faces. You will now see the front face of the cube remove to reveal the inside. If you now select the right wall four vertices and press the **W** key, you can select Subdivide to double the vertices. Repeat this procedure and add a number of vertices to the wall. You can now select vertices and remove them with **X** key to create a typical window. To finish off I have inserted a monkey in the scene by selecting Add>Mesh>Monkey from the Toolbox. You can see the completed scene in **Fig1**.

To add some atmosphere to this scene, we will make use of Volumetric Lighting. This type of lighting will show particles and dust in the light and give the impression of light flowing in from the light source. Although we have covered this in a previous issue, we have not applied this type of lighting through a window. In our current scene, we have a window with a cross in the middle to separate the four panes, and when volumetric light is applied to scenes such as this, you can get interesting effects. The reason why this particular type of window works well with volumetric light is because the cross in the middle is shadowed across the room, and gives the room a shady night look with moonlight flowing in.

The first thing to do is to add a spotlight to the right of the window so it shines in. You can do this by adding a light, increasing the size of the target area and position it so it is rotated down shining through the window. Imagine this is sun or moonlight and try to get a good angle on the light coming in; we

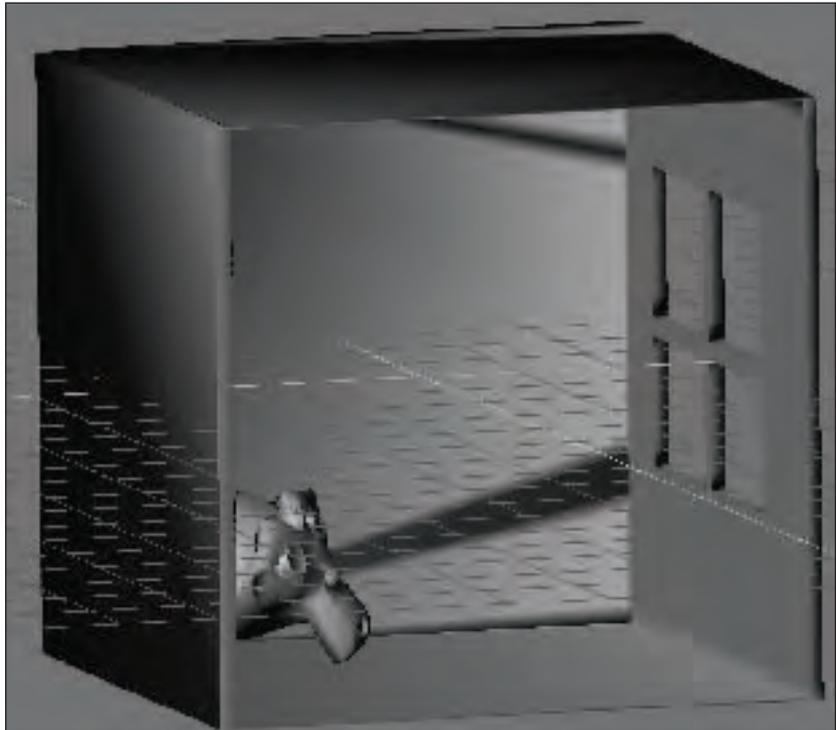
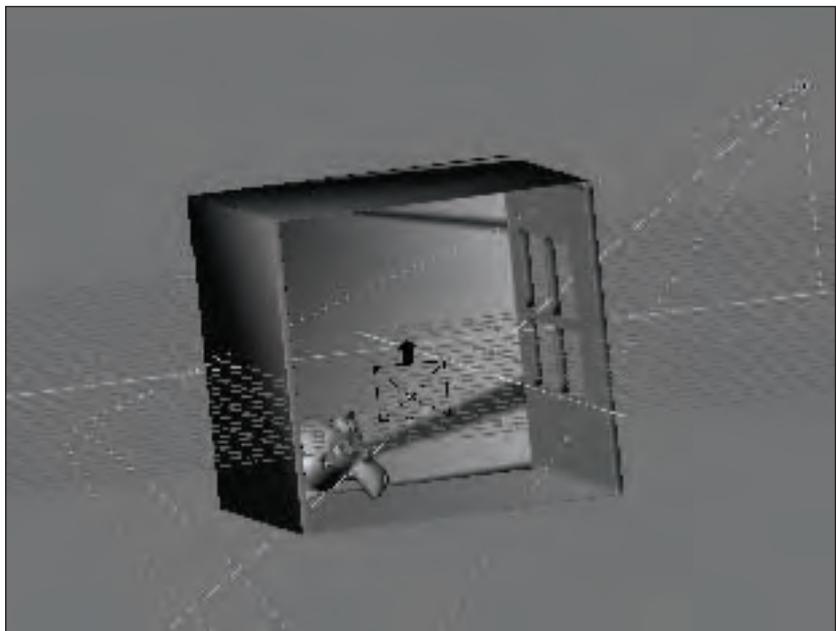


Fig1 A simple scene to begin the exercise of adding atmospheric lighting.



are aiming to get the light to flow into room. You can see the intended direction of the lamp in **Fig2**.

With the spotlight in place, we can now set the halo of the light. You can do this within the lamp buttons (**F4**), and there is the Dist (halo distance) and HaloInt (halo intensity) buttons that can be adjusted. Play with this buttons for a while, and you will want to keep adjusting the target size of the spotlight to get as realistic effect as possible. You can adjust the Energy slider to determine how bright the light is – this is often used to determine whether the scene is set in the day or at night. When you have played with these settings, you should have a scene fairly similar to that of **Fig3** overleaf.

Although our scene is looking interesting, we are essentially only seeing part of the room; the part where the light drops. In a **>>**

Fig2 A spotlight entering the room through the window can be moved and/or have its intensity varied to create differing atmosphere any time of day or night.



Fig3 A light shining in the room. Notice the projection of the window.



Fig4 An additional low power light shows the darkness of the room.

◀◀ realistic setting with this kind of light shining in, your eyes would have adjusted to the darkness of the room and you would be able to see the extent of the room in the darkness. To get this kind of added effect, we need to add another lamp at the front of the scene that has a low amount of light emitted out of it. When you have added the additional light, you may need to modify some of the settings of your main spotlight. Remember that light interacts with other light, and one setting may not necessarily apply when you add additional lights to the scene. You can see the result of this extra light in **Fig4**.

Radiosity

One of the problems with 3D modelling tools, are their limited application in calculating light effectively in all situations. There are various methods of positioning lights to artificially gain decent results, but the problem is more of an issue with the interaction of light on different objects. As an example, if you look around you now, you can see that light is emitting from various parts of the room and light reflected from one source is combined with

ATMOSPHERE TIPS AND TRICKS

A few handy hints

There are many different ways of adding atmosphere to your scenes, and many of them are simple and involve a few simple techniques. Here are some useful tips and tricks to get started with:

CAMERA ANGLES Camera positioning and angles can give a vastly different feel to a scene. As an example, if the camera is on the floor looking up, it gives a sense of height in the scene. If the camera is behind an object and part of the object is in view, you can create the effect of hiding in the scene. Remember that the camera is the viewers' eyes, so you need to be creative with it.

AMOUNT OF OBJECTS The number of objects can change how claustrophobic your scene feels. If there are many objects close together and there is no space between the camera and the object, the scene can feel small and limited. If you set the camera in the middle of a room with only a couple of small objects, the scene can seem much more spacious and light.

AUTHENTICITY If you are planning on creating a scene that you want to represent reality on earth, always ensure that your scene represents reality itself. As an example, in this type of scene, do not have objects hanging in the thin air or have humans picking up cars. In your animations ensure that physics are closely monitored and are modelled correctly.

the different properties of the different materials in the room. The basic premise of these events is that some incredibly complex light processes are going on, and 3D modellers need to replicate these processes as effectively as possible.

One interesting feature in *Blender* that can be used to solve some of these light processes, is called Radiosity. The basic idea is that each element in your scene acts as a receiver of light and you have a central emitter that distributes this light. This may sound just the same as adding a normal light to your scene as the emitter, but it really means setting a mesh as your emitter. This not only removes the need for lights in your scene, but allows you to control how your lights look and behave and how the light is emitted. As an example, you could create long neon tubes, thin light strips, circular lamps and more.

We will take a look at creating a simple scene with radiosity. The first thing you should do is create the same scene with the monkey head in the corner of the room. Ensure that you have removed all lights from the scene, and delete all materials that may be applied to objects. Next, add a plane to the scene and resize it so that it is a little smaller than the monkey head. You can now position it under the head so it looks as if the head is sitting on it. We will now light the room from this plane so it looks like the monkey head is sitting on the light.

To set the plane as an emitter, select it with the right mouse button and create a new material for it. Use the material buttons to set the colour of the light to be emitted, and then adjust the Emit slider up slightly. This control lets you set how much light will be emitted from the plane, and it is very sensitive. With the slider increased just slightly we can now begin the radiosity process. First, click on the Radiosity button (the button looks like a nuclear sign), and you will see a number of different buttons for controlling the scene.

We will only look at some of the essential buttons to get our radiosity working, but it is recommended that you fiddle around with these buttons and see their different effects.

NEXT MONTH

We will be continuing to consolidate more of the skills we have already touched upon earlier in the series...

If you've enjoyed following these *Blender* tutorials, make sure that you don't miss our six-page special and competition in *LXF50*, on sale Friday 23 January 2004.

To begin the process we need to select all of the objects in the scene by pressing the **A** key and then specify that these objects will be affected by the light from the emitter. You can set this by clicking on the Collect Meshes button. When you have clicked the button, you can begin the radiosity generation by clicking on the Go button. You will now see the scene build up the shading in each view and the light processes displayed. You can see this process as it works by seeing the wireframe guides where parts of the scene are being generated. There are three important buttons that affect the quality of the scene which are Wire, Solid and Gour and can be seen in the Radiosity Types box on the right.

Radiosity is an effective technique if you are making use of the *Blender* game engine. The benefit of the technique is that when you have created a radiosity effect that you are happy with, it can be saved and does not need to be recalculated each time. This is an important factor for interactive games, and you can create atmospherically lit scenes that are pre-processed for you. You can save your scene by clicking on the Add New Meshes button. This is not currently something that is undo-able, so ensure that you save your scene before you click this button just in case you change your mind.

Conclusion

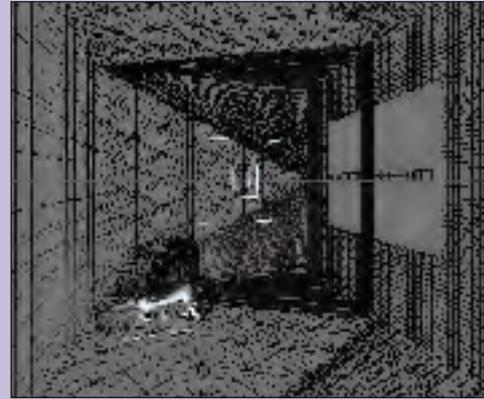
In this issue we have looked at some of the issues involved with creating atmosphere in your scenes with some examples of using lights effectively. It is recommended that you try many different light combinations to get different effects, and that you begin bringing in different techniques to your scenes such as particle systems to create realistic smoke inside your carefully lit scenes. Each of these techniques needs to be combined together to get the desired atmosphere, and it is always important to have in your mind the desired effect.

There are of course, many different techniques available for creating these effects, and it is wise to come up with different ideas for applying your knowledge. As an example, set yourself a project each time to create a happy scene, a scary scene, a cold scene, a wet scene etc. Using your knowledge, try to approach each scene to get the most out of its intention, such as creating cold, harsh objects for a cold scene. Practice makes perfect and the more you approach complex problems to model, the more you will encounter opportunities to add realism. 

RADIOSITY TYPES

Same scene, different effects

The three main types of radiosity supported in Blender are wireframe, solid and gouraud or 'gour' shaded types. (The latter type is sometimes also spelt 'gouraud'.) Each of these types has different effects, and if you are using radiosity in a *Blender* game, you may want to determine how much processing power is required for the high quality gouraud shading and if it is too high, use the lower quality solid type.



Wireframe shading.



Solid shading.



Gouraud shading.

TUTORIAL GIMP



CREATING PERSPECTIVE

Playing with shadows – adding depth to 2D images

They add style, depth, and in most cases they come with one click. Shadows are simple, fast, and a must-have for any digital artist, says **Michael J Hammel**.

As the acronym implies, *The GIMP* is an image manipulation tool. Its strength is in changing distinct pixels without knowledge of the image content. In the world of 3D, however, tools are required to understand what an object is and that it has depth. So how can we get *GIMP* to understand objects and depth? The most obvious way is with

shadows. *The GIMP* has many ways of dealing with shadows. Several come with built-in filters while others require minimal knowledge of basic *GIMP* features.

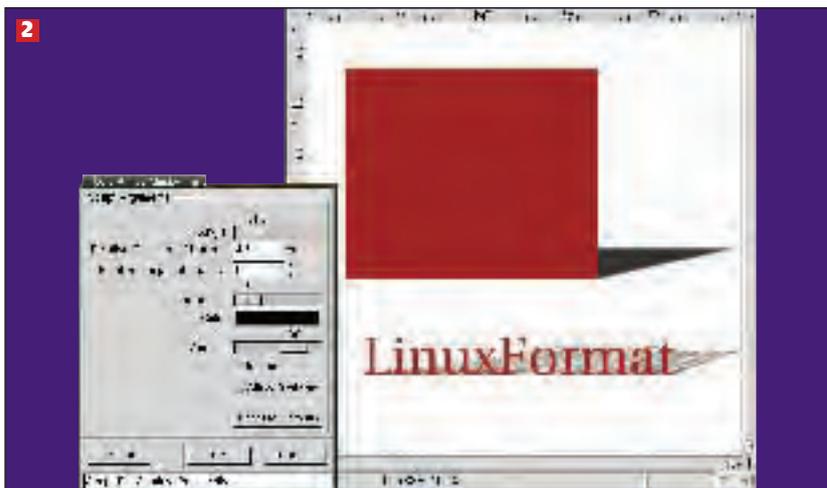
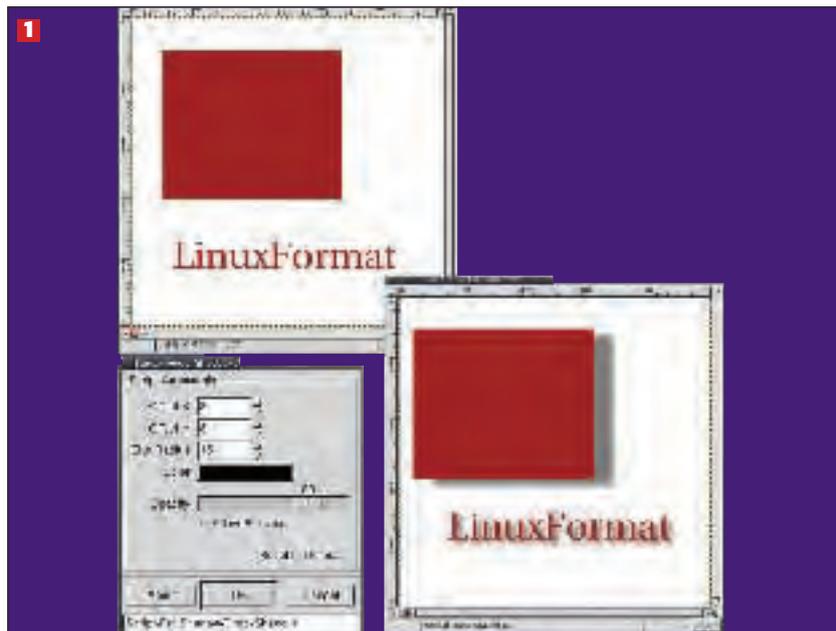
This month's tutorial takes quick peaks at *The GIMP's* built-in shadow tools, a method for generating your own perspective shadows, and a quick way to add a shadowy depth to any text.

STANDARD SHADOWS

The *GIMP* provides three standard shadow filters in the Script-FU>Shadow menu: Drop-Shadow, Perspective, and Xach Effect. For most situations a drop-shadow will suffice. For advanced artists, perspective shadows add flair by changing the perceived direction of incoming light. In any case, a quick look at all three menu options is all that is needed to learn their secrets.

Drop Shadow

1 A drop-shadow is simply a shadow placed behind an object. In this example we've drawn a red box and some red text, each on their own layer with a white background layer below them. Each layer is made active and the Script-FU>Shadow>Drop Shadow option is chosen. This filter provides options for setting how dark the shadow will be (its opacity) and how far offset it should be. Larger offsets and larger blur radii can make the shadow appear further back from the object. The effect is to make the object appear as though it floats above some surface.

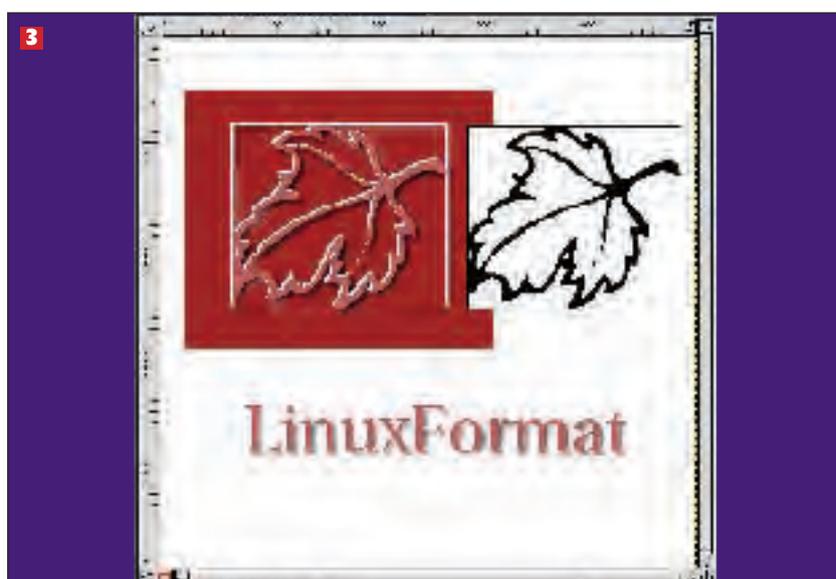


Perspective Shadow

2 An alternative to a drop shadow is a perspective shadow. A perspective shadow is like your own shadow when the Sun shines – it can run in front of you on the ground or behind you depending on the direction you face. In this example, the red block and text appear to have the Sun in front because the shadows run behind them. Because the shadow touches the bottom of the objects the effect is to make the object appear to sit on a surface. The Script-FU>Shadow>Perspective filter is used to create these shadows. The angle runs counter-clockwise starting at 0 along the X axis, so lower values create perspectives that are more flat.

Xach Effect

3 Also included as a shadow option is the Xach Effect filter. While this filter doesn't actually produce a shadow as its primary effect, it does include a shadow. To use this filter, make a selection over your object and apply the filter. The effect makes the selected area of the object appear to be raised and slightly translucent. In this example a leaf image with a black border was used to create a selection and then applied to the red box. The effect is slightly different when applied to a selection made from the text because the selection matched the shape of the object (the text). The result is a variation on a 3D text effect.

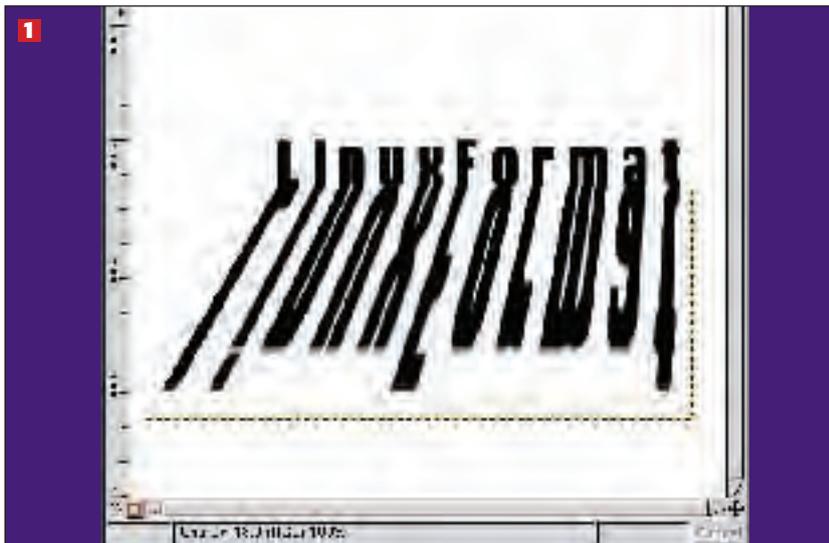


« PERSPECTIVE SHADOWS

The standard Perspective shadow filter provides plenty of configurable options, but it lacks an interactive feel. Adding perspective shadows manually is rather easy and gives you more control of the effect.

Manual perspective for text

1 Duplicate a text layer and flip the duplicate vertically. Then align the bottom of the original layer with the top of the flipped copy, essentially aligning the baselines of the text for each layer. Using the Transform tool from the Toolbox, scale the duplicate layers height by dragging down the bottom edge of the transform outline. Don't stretch its width, just its height. Apply this change, then switch to the Perspective mode of the Transform tool and drag the lower left corner out to the left. Apply this change.

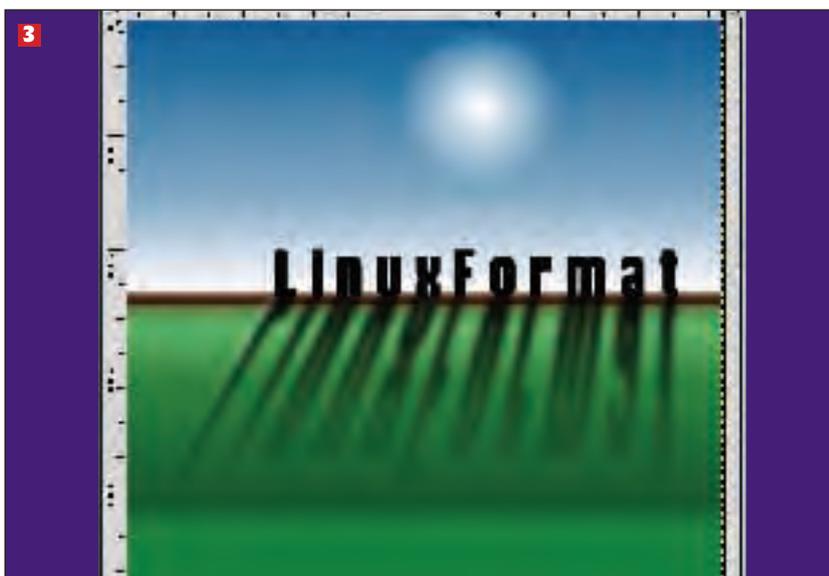


Blurred duplicate

2 Change the layer boundary size for the transformed layer to add some space to the bottom and right sides. You don't want the top to have any additional space and you don't really need any additional space on the left. The extra space is for blurring – be certain the layer has Keep Transparency turned off. Now Gaussian Blur this layer. This example used 10 pixels, but you can use whatever looks appropriate. Notice that after the blurring you won't have any blur along the top edge – that's why you didn't want to add any space to the top of the resized layer! That edge needs to match evenly with the original text. A shadow is most distinct where it touches its object but fades as it moves away from that point.

Final image

3 Now add a layer mask to the transformed layer. Apply a linear gradient going from black to white, starting at the bottom of the layer and running to the top. You need to play with this a little to get the right effect – you want the shadow to appear to fade away with distance. Add a gradient background. Finally, add a radial gradient (white to black) in a new layer and set its layer blend mode to Screen. The result is an image with a light source behind your object (text in this case) and a perspective shadow flowing toward the viewer.

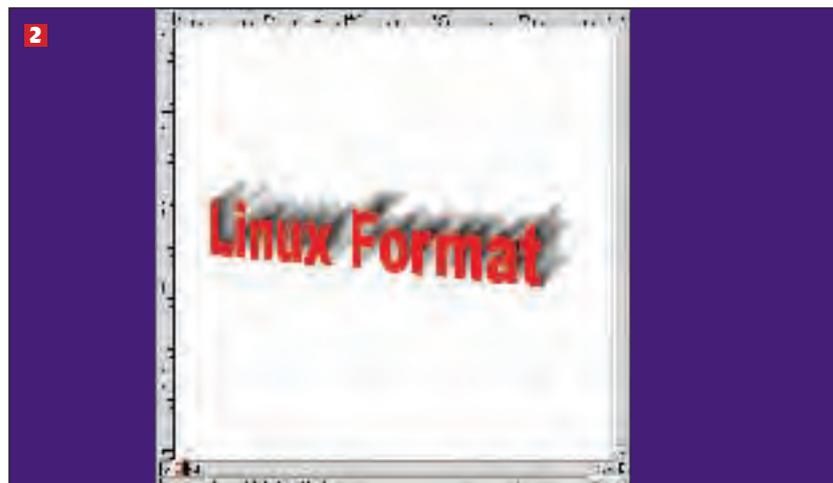


EXTRUDING TEXT

Shadows can also be used to simulate depth, providing a 3D effect to a 2D object. Although *The GIMP* does not provide a filter to extrude text, a blurred image can be used to simulate depth. The object will appear lit from in front and shadowed along its depth. This short tutorial takes only a few minutes to perform but can add plenty of pizzazz to any image. Note that this method will work best with thick, solid shaped fonts. It does not work well with script, thin-lined or Dingbat fonts.

Coloured text transformed

1 Start with a white background layer and add a coloured text layer. This tutorial only works if your original text colour and the extruded shadow are different colours. In this example, coloured text has been added to its own layer, sheared with the Transform Tool and the perspective changed (also with the Transform Tool). The layer border has been extended (see Layer Boundary Size in the Layers menu) to make space for the extruded area we're about to create. Duplicate the text layer. Use Alpha to Selection on it and fill the selection with black (or some colour other than the colour of the original text). Turn off the selection (Shift-Ctrl-A).

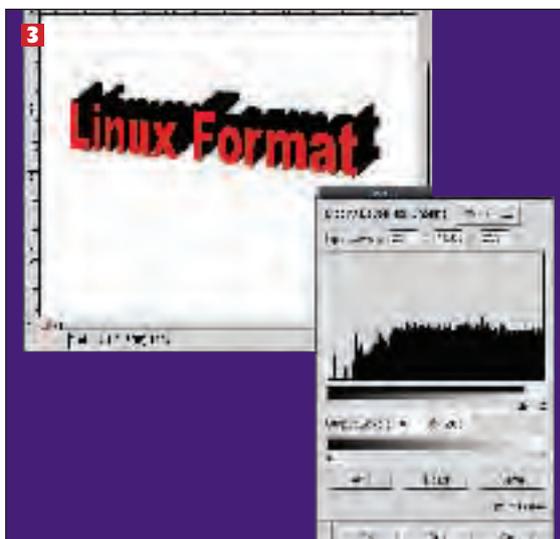


Extrusion started

2 Turn off Keep Transparency for the duplicate text layer and apply a motion blur (Filters > Blur > Motion Blur). The length of the blur shouldn't be too large - you don't want really deep 3D extrusion. Here a length of 35 pixels was chosen for a 512x512 canvas. The angle is set with 0 (zero) being to the left along the X-axis, so if you want extrusion to the upper right (as in this example) subtract the angle desired (40 degrees using the Measure Tool to estimate) from 180. This example used 140 degrees to make the motion move from lower left to upper right, which matches the transforms performed on the original text. Use the Linear option for the motion blur. Note that if you don't have enough border or you use too long a blur length or you use the wrong angle the extrusion will get chopped off at the layer border.

Extrusion completed

3 Move the duplicate layer below the original text layer and then merge it with the white background. You must do this or the next step won't work right. To make the blurred region appear solid, use the Levels dialog and drag the black and gray handles in the upper adjustment bar to the right. This will make most of the pixels extremely black. If you drag the black handle all the way to the right the edges of the text become jagged, so you may want to leave a little space between the white and black handles, with the gray handle set somewhere between or on either end.



NEXT MONTH

We've mentioned the subject of Filters before when considering non-destructive editing - next time we concentrate on them in more depth.

PROBLEM-SOLVING

Practical PHP Programming

On the fourth day of Christmas, **Paul Hudson** gives you; four calling birds, three French hens, two turtle doves, and a partridge in a PEAR tree...

Writing PHP code is quite a fun task – as we're sure of you will agree – it's an easily learnt, flexible, and feature-rich language that lends itself very well to scripts of all shapes and sizes. Almost certainly as a result of this comparative ease of use, many PHP programmers take to re-inventing the wheel a lot and spend much time writing and rewriting 'new and improved' code to handle authentication, caching, forums, opinion polls, and other systems that have been coded and recoded in the past.

This is not unique to PHP as a language. C++, for instance, utilises the Standard Template Library (STL) to provide abstract data types (ADTs) for programmers that are pre-written, pre-optimised, and are cross-platform. As a result, a programmer writing for Windows using C++ can create and use STL stacks, vectors, maps, and queues, and have them recompile smoothly on Mac OS, OS/390, Solaris, and, of course Linux. This is immensely time-saving, so it's unsurprising that many other languages have their own versions of this plan. In PHP, this is known as PEAR, the

PHP Extension and Application Repository, and is available from <http://pear.php.net>. The purpose of this tutorial is to briefly outline what PEAR can offer, with a particular focus (thanks to popular demand) on PEAR::DB, the advanced database abstraction library that forms the most popular part of PEAR.

All prePEARed?

Switching to an off-the-shelf solution for a variety of common problems might seem like a hard task – after all, each problem is unique, right? In most cases, yes – we want to solve a problem that has yet to be solved. This is what makes programming a fun and challenging task, and, quite frankly, the problem wouldn't be a problem if it were already solved! However, while most problems are unique, they can usually be broken down into a set of smaller problems, of which each can be broken down into a smaller set of problems, and so on. This recursivity was first put into poetry by Jonathan Swift, and latterly updated and popularised by famed Victorian mathematician Augustus De Morgan:

**Great fleas have little fleas
upon their backs to bite 'em,
and little fleas have lesser fleas,
and so, ad infinitum**

Of course, problems (and fleas) can only be reduced to a certain level, at which point you actually need to go about solving them. However, it's at this point – when we've broken our large, complex, and unique problem down into itty-bitty parts – that we find we can re-use components easily, and can indeed benefit from doing so. For example, handling file uploads is a fairly common task that is usually one small part of a larger system, and there are many ways you can handle it. Now, imagine for a moment you're joining a large company where file uploads are handled, and you ask what method was used to handle the files – would you rather hear “it's our own custom code, written by one of our junior programmers”, or “we use the PEAR code”?

Unless you enjoy suffering, you'll certainly prefer the option that gives you code that's written by a set of experts in their field, tested by thousands, written according to strict style rules to make the source code easy to read, and also designed to be flexible. This is where PEAR wins out, and once you make it over the initial barrier to entry – that is, finding one or two places where PEAR can help you immediately and give it a try – you'll almost certainly stick with it.

If you really get into the swing of using PEAR, you may even find it worthwhile to build up your own modules that are compilations of PEAR modules, which models your code along the same building-block lines as PEAR itself. If you manage to achieve this nirvana – it helps if you use OOP – then you'll find you can turn around even the largest projects by putting together your small building blocks into several larger blocks, then those larger blocks into one finished solution, thus completing the picture outlined by the poem's little-known second stanza:

**And the great fleas themselves, in turn,
have greater fleas to go on;
While these again have greater still,
and greater still, and so on.**

ComPEARing the packages

While most of the requests we've received to cover PEAR have centred on PEAR::DB, we can hardly cover PEAR and not discuss the dozens of other modules that make it up. So, before we jump into the ins and outs of PEAR::DB, let's first look at what else PEAR can do to revolutionise your PHP programming:

ACCELERATE YOUR CODE

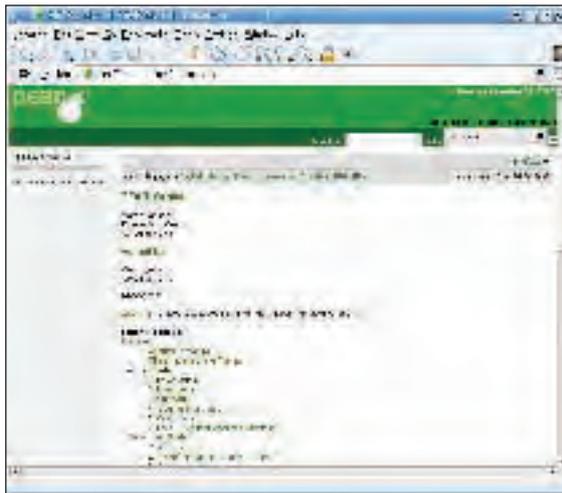
Back in *LXF34* we reviewed various PHP Accelerators to see which would provide maximum acceleration for your PHP scripts, and at the bottom of that test, with a lowly 3/10, came the *Alternative PHP Cache*. Since then, APC was contributed to the PEAR project and taken onboard by George Schlossnagle (the author of *Advanced PHP Programming*, ISBN 0-6723-2561-6), where it has been revolutionised into a much faster and leaner product that competes well with *Zend Performance Suite*.

BENCHMARK PERFORMANCE

If you're sick of trying to hack together a decent profiler with `microtime()`, PEAR's benchmarking framework is just what you're looking for – you can check and time performance data based upon overall script execution, as well as more specific things such as individual function execution.

PARSE XML

Although PHP's XML-parsing featureset is powerful and will be



The PEAR manual (online at <http://pear.php.net>) is modelled after the main PHP manual, but sadly it doesn't manage to be nearly as comprehensive – why not contribute a little time to help out?

even better in PHP 5, PEAR contains a number of packages that provide more specific XML parsing capabilities such as RSS and RDDL handling.

FIND FILES

PHP has solid filesystem-handling support, but, as with XML, the PEAR developers have found ways to make things even easier for us. For example, finding files in a given directory is now as simple as using the pre-written directory search component bundled with PEAR.

While most parts of PEAR are easy as easy to use as any other part of PHP, the most popular and most advanced part of it also happens to be the most complicated to fully grasp. Having said that, PEAR::DB, for it is that of which we speak, is worth spending a few hours tinkering with, as it's actually rather good.

Don't disPEAR!

Although it's hard to learn fully, we're going to look over some of the various features of PEAR::DB in a simple way – hopefully everyone should be able to see why it's a good thing. First, let's take a look at what PEAR::DB offers:

- Complete database abstraction for a variety of databases (see the box *PEAR::DB DBs* on the following page for a list)
- Object-orientated querying
- A selection of quick 'power' functions that perform a query and extract data in one call
- Prepared statements

We'll be looking at each of these, so if you're not sure what a prepared statement is you can relax as it will be covered shortly. The first two items there, though, are fairly elementary, so we can just dive right in with a lump of code:

```
<?php
1 include_once('DB.php');
2 $conn = DB::connect("mysql://user:pass@server/mydb");
3 if (DB::isError($conn)) {
4     echo 'Connection error: ', $conn->getMessage();
5     exit;
6 } else {
7     $result = $conn->query("SELECT Field1, Field2 FROM
8     sometable;");
9     while ($result->fetchInto($row,
10     DB_FETCHMODE_ASSOC) {
11         extract($row);
12         echo "$Field1 $Field2<BR />";
13     }
14 }
```

DOG FOOD

It's a commonly held view that a company (or in our case, a magazine) should eat its own dog food, by which it is meant that a company should follow its own rules and internally use any software it produces. If you read that re-inventing the wheel is bad and thought, “Hey, wait a minute – LXF is redoing its website and creating its own forum system! Why didn't they use an off-the-shelf system?!” then we're afraid it's simply a case of “do as we say, not as we do”. Anyway, our wheel is going to be so much better...

TUTORIAL PHP

```

<< 12 $result->free();
    13 }
    14 $conn->disconnect();
    ?>

```

That's quite a bit of code to absorb all at once, so let's look at the important lines...

LINE 1 This includes the PEAR::DB code that's necessary to connect and query using PEAR::DB. Note that we just specify **DB.php**, as the PEAR include directory is usually in the PHP include directory directive by default.

LINE 2 This connects to your database using username **user**, password **pass** on server **server**, then selects the database **mydb**. The **mysql://** part specifies that this is a MySQL server, and you can replace that with one of the other databases listed in the *PEAR::DB DBs* box below to change the connection type. Note that **DB::connect()** returns a database connection object, which is stored in **\$conn**.

LINE 3 This will return true if the last function called on **\$conn** returned an error.

LINE 4 **getMessage()** returns the last human-readable message stored in the connection. In the case of an error, this will store a small amount of text describing the error.

LINE 7 **query()** simply sends a query to the database object and returns the result as an object. To maintain the cross-platform advantages of PEAR::DB, steer clear of non-ANSI SQL extensions, such as the LIMIT clause in MySQL. Note that DELETE, INSERT, and UPDATE queries do not return an object because they don't return any data – instead you can expect true (**DB_OK**) or an error.

LINE 8 **fetchInto()** is the equivalent of **mysql_fetch_array()**, and takes one row from the result, and places it as an array specified in parameter 1. The second parameter specifies what data is stored, and is optional – **DB_FETCHMODE_ASSOC** will make the first parameter into an associative array, with field name being used as the key.

LINE 12 This frees the result, which is always a smart move to keep your 'house' clean.

LINE 14 Finally, we disconnect from the database, again another good 'housekeeping' move.

Power functions

Point three of the four helpful things that PEAR::DB does for you is provide power functions that accomplish a lot of functionality in just one line of code. The simplest of these is **getOne()**, which takes a simple SQL query to execute on a database connection, executes the query, and returns the selected value. For example:

```

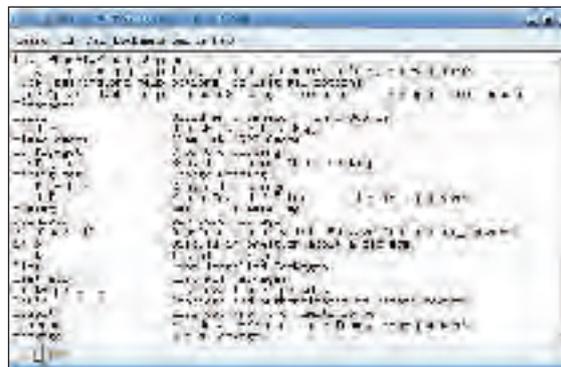
$averageage = $db->getOne("SELECT AVG(Age) FROM
userlist;");
echo "The average member age is $averageage<BR />";

```

Naturally this isn't much good for more complicated queries as it only returns the first field from the first row, however there are other calls that add a little more flexibility: **getRow()**, for example, executes the query passed and returns all columns of the first row returned by the query, whereas **getCol()** returns the first column of all rows in the query. Both **getRow()** and **getCol()** return their information in an array, whereas **getOne()**, as you've seen, returns a simple variable.

Query meta-information

As well as reading the actual results of your query, you can also read back various information about what the query returned –



If you installed PHP yourself, use the **pear** command from the shell prompt to search for and install PEAR modules.

the equivalent of **mysql_num_rows()**, **mysql_num_cols()**, and **mysql_affected_rows()** all have their counterparts in PEAR::DB in the form of **numRows()**, **numCols()**, and **affectedRows()**.

As with MySQL, **numRows()** and **numCols()** should be used when data has been selected and you want to know how much data there is, whereas **affectedRows()** should be used when data has been changed (*ie* you've used DELETE, UPDATE, or INSERT) and want to know how many rows have changed as a result of your query. In use, **numRows()** and **numCols()** work alike, whereas **affectedRows()** is slightly different. Put simply, **numRows()** and **numCols()** are called upon the data result itself and simply counts the number of items in the data set. On the other hand, **affectedRows()** is used for queries where no result is returned, and so you need to call it using the main database object itself.

To make this clear, take a look at the following example:

```

<?php
    $result = $db->query("SELECT User, Pass FROM
memberlist;");
    echo 'Rows returned: ' . $result->numRows() . ' rows\n';
    echo 'Cols returned: ' . $result->numCols() . ' cols\n';
    $result = $db->query("INSERT INTO memberlist VALUES
('Luke', 'Jedi');");
    echo 'Rows affected: ' . $db->affectedRows() . ' rows\n';
    $result->free();
}
$db->disconnect();
?>

```

The calls to **numRows()** and **numCols()** are both called on the **\$result** object returned by the SELECT query, but note that in the line indicated in red above, the return value from the INSERT query *isn't* used to call **affectedRows()**. This is because queries that change data rather than return data don't send back a resource object, so you need to go direct to the DB. Keep this in mind, because if you try to call **numRows()** or **numCols()** on the return value from a non-SELECT query you'll get an error.

PrePEARed statements

By now you should be able to see that switching to PEAR::DB means you can centralise all DB-specific code so that the only line that varies across databases is the original connection – this is a great benefit by itself, but there's a lot more that PEAR::DB can do to make itself stand out amongst the crowd even more.

Perhaps the most helpful feature is prepared statements, which allows you to define approximately what a query will do without having to pass exact values into it. For example, you can prepare a query that will insert rows into a field without specifying

PEAR::DB DBs

PEAR::DB supports the following database systems – use the first part as part of the connection string:

fbsql	FrontBase
ibase	InterBase
ifx	Informix
msql	Mini SQL
mssql	Microsoft SQL Server
mysql	MySQL
oci8	Oracle 7/8/8i
odbc	ODBC (Open Database Connectivity)
pgsql	PostgreSQL
sybase	SyBase

what the field values will be, then, later, you can call that query with your values.

This might all sound complicated at first, but essentially it makes an SQL query almost like a function – what your prepared statement does won't change, but the values it works with will. If you don't think this is helpful, consider how hard it is to read long SQL queries – if you're specifying twenty fields and values it can get complicated to understand precisely what's going on! All this changes with prepared statements, as, because the prepared statement already knows what fields it's working with, you only need to specify the values; you essentially split the command from the data.

Your basic prepared statement looks like this:

```
INSERT INTO memberlist VALUES (?, ?);
```

Each question mark there symbolises one field for which we don't currently know the value, and you can have as many question marks as you have fields. You can also mix and match question marks and hard-coded values, like this:

```
INSERT INTO memberlist VALUES (?, 'blue', ?, 10, 0, ?, ?);
```

Here the second, fourth, and fifth values are hard-coded as 'blue', 10, and 0, but the first, second, sixth, and seventh values are variable and need to be provided later. Note that the question marks need to be exactly as shown, as anything else is considered to be a hard-coded value. For example, if you want all rows in the first field to start with 'foo-', you couldn't use a query like this one:

```
INSERT INTO memberlist VALUES ('foo-?', 'blue', ?, 10, 0, ?, ?);
```

While it might seem logical that the question mark in 'foo-?' is treated as a wildcard and replaced with the value later on, it isn't, so repeat after me: anything other than a plain question mark is treated as a hard-coded value.

With that in mind, let's take a look at a working code example of prepared statements. For this example, you'll need a table, in this instance "future_publishing", that has three fields: "user", "pass", and "magazine".

```
$teamlxf = array(
    array("Matt", 'peekaboo'),
    array("Nick", 'tulip'),
    array("Julian", 'pembrokeshire')
);
$stmt = $db->prepare("INSERT INTO future_publishing
VALUES (?, ?, 'LXF');");
foreach($teamlxf as $teammember) {
    $db->execute($stmt, $teammember);
}
```

You'll need to wrap all that up in the connect/disconnect code from earlier before you try it. The two key parts there are **prepare()** and **execute()**, which take two parameters and one parameter respectively. The **prepare()** function is used, unsurprisingly, to prepare a statement, and it takes the statement you want prepared as its parameter, returning the number that's been assigned to the prepared statement. In our code this value is kept locked away in **\$statement** for use two lines later.

A **foreach()** loop is used to cycle through each team member in the **\$teamlxf** array, which is an array of arrays. Each element in **\$teamlxf** is an array containing a username and password for our prepared statement, and the **foreach()** loop extracts each team member and sends it off to the **execute()** function as parameter two. So, parameter two for **execute()** is the data we want to send in; parameter one, as you can see, is the prepared statement to use. Note that the question marks in the prepared statement are filled with the **\$teammember** array from

\$teamlxf directly – you don't need to break the array apart to send them off to **PEAR::DB**, which makes your code a great deal easier to read.

ImPEARed performance?

Common sense says that directly calling PHP's MySQL functions will be faster than using an abstracted code library written in PHP. After all, every call you make using **PEAR::DB** needs to jump through a variety of hoops to handle the variety of databases, whereas calling **mysql_query()** is entirely unambiguous and therefore likely to be much faster.

In practice, **PEAR::DB** usually runs at about 1/2 the speed of the equivalent DBMS-specific code, although this may depend on the kinds of query you're running. Of course, it's impossible for **PEAR::DB** to ever outpace DBMS-specific calls simply because once the database abstraction has been resolved internally to **PEAR::DB** it will go ahead and execute the native calls itself.

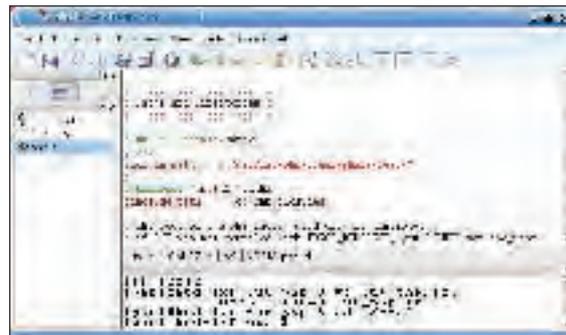
However, you need to weigh this hefty performance hit against the extra flexibility offered by **PEAR::DB** – if you sprinkle **mysql_*** calls throughout your script then later want to port to Oracle, you'll find yourself with a fair amount of work on your hands. If you use **PEAR::DB**, you just need to change the connection line, which is a darn sight easier. The addition of prepared statements is also quite attractive, particularly when you need to handle large amounts of fairly similar data, but also because *ASP* programmers (using Microsoft's ADO technology) have had prepared statements for some time, and so it makes the transition to PHP easier for them.

Where's the partridge?

PEAR, particularly **PEAR::DB**, offers an immense amount of code that's easily available and, more importantly, easy to re-use across projects. While we've focused on **PEAR::DB** here, simply because it's the most advanced module in **PEAR**, there's a lot more out there – did you know, for example, that you can grab a *Mono* (yes, the cross-platform .NET system by Miguel *et al*) extension for PHP? No? You should check the **PEAR** site more often!

So, in conclusion, **PEAR** is a great resource for PHP developers who want to take their projects a little further – either by using the advanced components on offer there, or by using the basic components there and thereby spend their programming time working on the more interesting and challenging topics. Even if you're not sure whether you want to commit to using **PEAR** just yet, at least give it a look over and see whether it has anything that can directly help you.

Alright, so we lied about giving you a partridge... **LXF**

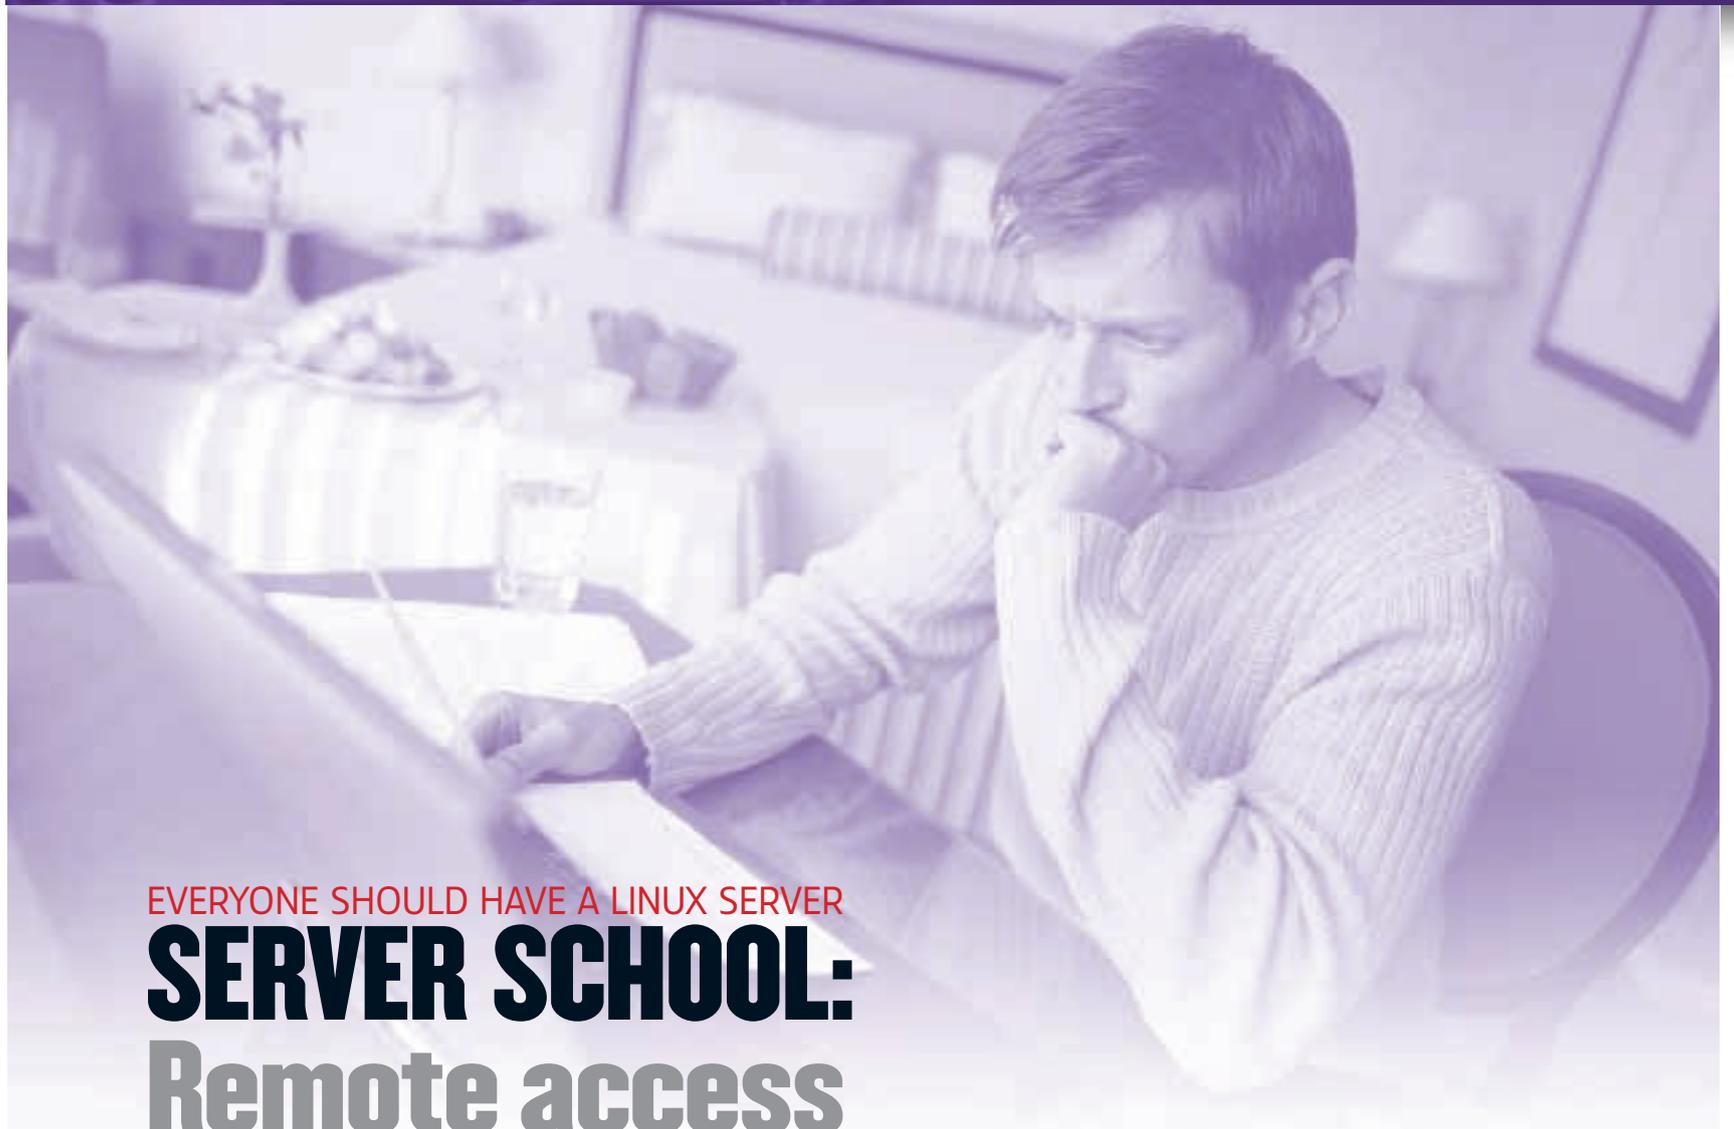


To use **PEAR** files easily within your scripts, be sure to modify your **php.ini** file so that the **PEAR** installation directory is an include directory.

NEXT MONTH

In *Linux Format* issue 49, we'll be putting our noses to the grindstone with our first proper coverage of PHP 5.0.

As always, we're looking out for ideas for new *Linux Format* PHP tutorials. Send an email to me at paul.hudson@futurenet.co.uk with your ideas – all the good suggestions that you send in will be covered in future issues. If you're short of ideas, you're certainly welcome to write in or post on the forums at our website www.linuxformat.co.uk with comments – we're passionate about improving the overall quality of our tutorials!



EVERYONE SHOULD HAVE A LINUX SERVER

SERVER SCHOOL:

Remote access

PART 4 Keeping our servers close is not always ideal. David Coulson looks at how to maintain those that have packed their bags, bid a tearful goodbye and left home.

Since the conception of the Internet, the general consensus when provisioning bandwidth for a server is 'Take the server to the bandwidth, not the bandwidth to the server.' Unfortunately, since most of us don't have DS-3 or Ethernet circuits to Tier 1 providers in our office or house, it means taking the server off site to a collocation facility or data center to take advantage of the extensive and often cheap bandwidth available. This can often mean many miles between the administrator and the server, making on-site maintenance impractical, unless in an emergency situation where hardware needs to be changed or we actually lose connectivity to the device.

Notwithstanding major network failures, Linux is a wonderful system that allows us to remotely maintain a device without a great deal of effort. Whether it's a single web server, or a large cluster of systems providing application services to customers, there is very little that you can't do remotely to a running Linux system. Depending upon one's level of confidence in both the hardware and your own abilities, everything from upgrading the kernel to installing a new package can be done from thousands of miles away.

Telnet

Remote access at its most basic, and rudimentary, comes in the form of telnet, which allows us to bring up a shell prompt on a

remote system over TCP. The shell doesn't care if we're remote or local, so anything we would normally do when connected to the device via a standard keyboard and monitor, or over a serial link, can be done via a telnet session. Telnet is particularly dumb and does little more than read and write characters from a TCP session, making it a rather insecure way to do things. Passwords, usernames and indeed everything you type, fly back and forth over the TCP



There are cross-platform commercial admin tools available, like *NetOp*, but you can rely on free software too!

connection without any encryption. Anyone with a packet sniffer and a little creativity can come along and grab the TCP packets from the connection, causing either 'man in the middle' attacks by feeding the remote endpoint incorrect data to exploit it, or simply logging the session and using the username and password we used to get into the box to compromise it at a later date.

From a maintenance point of view, telnet is next to useless, but users should not discount it, as it is a particularly effective and flexible debugging tool for any service that runs over TCP. We can check our mail server to find out if it's an open relay using telnet, by connecting to it on port 25 and sending it some commands:

```
$ telnet 10.1.1.5 25
Trying 10.1.1.5...
Connected to 10.1.1.5.
Escape character is '^]'.
220 mail.mx.davidcoulson.net ESMTP qmail 1.03-djc-2
MAIL FROM: david@davidcoulson.net
250 ok
RCPT TO: bob@example.net
250 ok
```

In this case, we're connecting to a local system which permits relaying from the internal network, however doing external tests on a remote host should yield a very different response – if not, the administrator needs to jump and fix his mail server rather quickly to stem the flow of spam which is likely passing through it at an amazing rate. We can also use telnet to quickly verify what version of SSH we're running to ensure we've got the latest revision of the packages, ensuring our system is secured.

```
$ telnet 10.1.1.133 22
Trying 10.1.1.133...
Connected to 10.1.1.133.
Escape character is '^]'.
SSH-2.0-OpenSSH_3.6.1p2 Debian 1:3.6.1p2-9
```

We need a secure alternative to telnet to comfortably administrate a system remotely without having to worry about our packets being sniffed, which is where SSH comes into play.

SSSSH!

SSH is an alternative to *rsh*, which is a very simple and insecure remote shell utility, rather than a direct replacement to telnet. *rsh* was designed to allow users and administrators access to remote systems, generally within the same LAN or WAN, without having to authenticate using their standard UNIX password. Most of the authentication was done via IP-based permit lists, however the actual shell session is still wide open making it less than ideal. SSH is a great replacement for this tool, offering secure sessions as well as key based authentication, making it the more secure way to access a shell on a remote box. As SSH is simply a protocol, there are a few different implementations of it, although now that SSH2 is only freely available for non-commercial use, everyone is moving over to OpenSSH from the OpenBSD project. OpenSSH offers SSHv1 and SSHv2 protocol support, making it accessible using any client from commercial SSH, SecureCRT, Putty and MindTerm.

SSH is very simple to get up and running, as the vast majority of distributions come ready to roll with SSH. Simply point your client at the box, give it the username and password, and off you go. Rest assured that the username and password are sent to the other end of the connection wrapped in a secure encrypted protocol, making it somewhat difficult for someone to find out what authentication information is being used at the remote end



without a rather large supercomputer and oodles of times on their hands. One would hope that by the time they crack the encryption algorithm, generally taking many tens of thousands of CPU hours, the information they yield will be next to useless, or at least significantly less valuable than when they set out on their task.

SSH offers us far more than simply remote shell access. By authenticating using keys, we can shell into another box without supplying a password, making it ideal for running scripts from a cron job, or simply as a time saving exercise. Using public key authentication, we can safely throw around our SSH public key, placing it on remote boxes we need to have access to, but keeping our private key safely on our local system, or even on a USB flash device to ensure it is not stored on any system.

```
david@tailtiu:~$ cat .ssh/identity.pub
1024 35
14618708027357125607983662534277043338916017755
7852473731257385961486565457729119934917522699217
512274921919746199888128338788580252894681872745
0727814525135333112706786872339152974102310693916
576729167705554174279180919835964433629737828207
26101417912430853149908741215562357376211741326561
324000287066068057 david@niamh
```

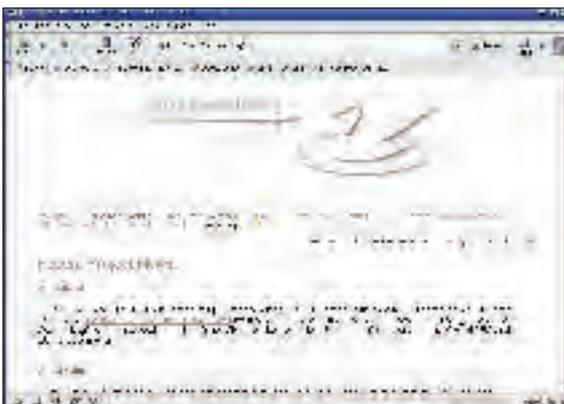
```
david@tailtiu:~$ ssh -v 10.1.1.7
OpenSSH_3.4p1, SSH protocols 1.5/2.0, OpenSSL
Ox0090581f
debug1: Reading configuration data /usr/local/etc/ssh_config
debug1: Applying options for *
debug1: Rhosts Authentication disabled, originating port will
not be trusted.
debug1: ssh_connect: needpriv 0
debug1: Connecting to 10.1.1.7 [10.1.1.7] port 22.
debug1: Connection established.
debug1: identity file /home/david/.ssh/identity type 0
debug1: identity file /home/david/.ssh/id_dsa type 2
debug1: Remote protocol version 2.0, remote software version
OpenSSH_3.6.1p2 Debian 1:3.6.1p2-9
debug1: match: OpenSSH_3.6.1p2 Debian 1:3.6.1p2-9 pat
OpenSSH*
Enabling compatibility mode for protocol 2.0
```

SSH is a secure alternative to rsh and telnet, although it offers far more features than simple shell access.

```

<< debug1: Local version string SSH-2.0-OpenSSH_3.4p1
debug1: SSH2_MSG_KEXINIT sent
debug1: SSH2_MSG_KEXINIT received
debug1: kex: server->client aes128-cbc hmac-md5 zlib
debug1: kex: client->server aes128-cbc hmac-md5 zlib
debug1: SSH2_MSG_KEX_DH_GEX_REQUEST sent
debug1: expecting SSH2_MSG_KEX_DH_GEX_GROUP
debug1: dh_gen_key: priv key bits set: 134/256
debug1: bits set: 1585/3191
debug1: SSH2_MSG_KEX_DH_GEX_INIT sent
debug1: expecting SSH2_MSG_KEX_DH_GEX_REPLY
debug1: Host '10.11.7' is known and matches the RSA host key.
debug1: Found key in /home/david/.ssh/known_hosts:37
debug1: bits set: 1589/3191
debug1: ssh_rsa_verify: signature correct
debug1: kex_derive_keys
debug1: newkeys: mode 1
debug1: Enabling compression at level 6.
debug1: SSH2_MSG_NEWKEYS sent
debug1: waiting for SSH2_MSG_NEWKEYS
debug1: newkeys: mode 0
debug1: SSH2_MSG_NEWKEYS received
debug1: done: ssh_kex2.
debug1: send SSH2_MSG_SERVICE_REQUEST
debug1: service_accept: ssh-userauth
debug1: got SSH2_MSG_SERVICE_ACCEPT
debug1: authentications that can continue:
publickey,password,keyboard-interactive
debug1: next auth method to try is publickey
debug1: try pubkey: /home/david/.ssh/id_dsa
debug1: input_userauth_pk_ok: pkalg ssh-dss blen 434 lastkey
0x80faf18 hint 1
debug1: read PEM private key done: type DSA
debug1: ssh-userauth2 successful: method publickey
debug1: channel 0: new [client-session]
debug1: send channel open 0
debug1: Entering interactive session.
debug1: ssh_session2_setup: id 0
debug1: channel request 0: pty-req
debug1: channel request 0: shell
debug1: fd 3 setting TCP_NODELAY
debug1: channel 0: open confirm rwindow 0 rmax 32768
Linux maeve 2.4.20-pre10-fairsched #5 SMP Tue May 6
01:34:19 BST 2003 i686 GNU/Linux
Last login: Wed Nov 5 19:57:45 2003 from

```



FreeSwan provides an implementation of the popular IPSec tunneling protocol, included in latest kernels.

```
tailtiu.dmoz davidcoulson.net
```

```
david@maeve:~$
```

We can also do file transfer securely over SSH, using the rather nifty little *scp* utility. Dumping data from one system to another is trivial using *scp*, as the syntax is very much like the standard *cp* tool, although one needs to pass a hostname for it to connect to in order to copy the file. As SSH includes compression support in the protocol, if we're copying tar files or plain text, it will be compressed on it's way through the network, improving throughput at the cost of CPU utilisation. We can also tunnel other transfer mechanisms over SSH, including *rsync* and *CVS*, making them significantly more secure than their standard implementations, as well as allowing them to make use of our SSH authentication methods for improved ease of use.

```

david@shuttle:~$ RSYNC_RSH=ssh rsync -aWvc patch.diff
10.11.5:
building file list ... done
patch.diff
wrote 129 bytes read 36 bytes 110.00 bytes/sec
total size is 0 speedup is 0.00

```

Tunneling

While SSH is a great way to hit a single box, if we have more than one system, with one at the front and a whole load behind, SSH can become quite annoying. Shell into one box then shell into something else. Copying files is a nightmare and we can't cleanly test TCP connections on a remote system from our local box. The simplest way to handle this is by running an IP over IP tunnel, which essentially puts our local box onto the remote network. If the backend systems are running on a 10.0.0.0/8 address, or one of the many other private IP ranges, then our local system will end up with an IP from the same subnet so we can directly access systems sitting behind the firewall.

There are many IP over IP implementations, some are secure and can be considered rock solid for production use, and others are not quite so great. The most basic is the 'General Routing Encapsulation' tunnel, or GRE, which basically wraps the packets in another IP packet. IP-IP is also a similar implementation of an IP over IP service, and both of these require kernel modules to be loaded and the *iproute2* utility to be used to create the tunnel and setup the routing. As both of these don't securely encapsulate the traffic, we're not going to look at them too closely, although the Linux Advanced Routing and Traffic Control HOW-TO at <http://lartc.org/> has some detailed information for bringing up such a tunnel.

As an alternative we can either use the ever popular *IPSec* protocol, which is available for Linux 2.4 as part of the FreeSwan project at <http://freeswan.org/>. If you're crazy enough to run Linux 2.6 on something, then there is an implementation of *IPSec* ready to roll, although there is also a FreeSwan port to 2.4 for those of us who don't want to trust 2.6 with our lives. *IPSec* requires a kernel patch and rebuild on each end, making it less than ideal if your remote server is far enough away so that if the kernel decides not to work quite right, you're going to be stuck. On the plus side, *IPSec* can run tunnels to pretty much any device out there, from a Linux workstation to a Cisco router or many other firewalls or routing devices. Configuring *IPSec* is far out of the scope of this tutorial, although it has been covered in previous *LXF* issues in some depth.

We also have a variety of proprietary tunneling protocols for Linux, including *vtun* and *OpenVPN*. *vtun* has taking a beating

over the years due to its rather poor security implementation, although no one has yet come up for an exploit for it. It did bring us the TUN/TAP kernel module, which everything from *OpenVPN* to User-mode Linux uses for network connectivity, plus it is obscenely easy to setup and get running over either UDP or TCP. *vtun* is available for download at <http://vtun.sf.net/>.

vtun has a very basic configuration format that is almost identical on both ends, other than we change the IPs:

```
options {
port 12345;
ppp /usr/sbin/pppd;
ifconfig /sbin/ifconfig;
route /sbin/route;
firewall /sbin/ipchains;
}

default {
comp lzo:9;
encr yes;
keep yes;
mult killold;
type ether;
prot tcp;
persist yes;
}

tun-3 {
pass password;
up {
ifconfig "%% 10.2.1.3 netmask 255.255.255.0 mtu 1460
up";
};
}
```

OpenVPN on the other hand was designed to be extremely secure and offers more security features than most people can possibly comprehend. While it is a proprietary protocol, making it less ideal than *IPSec* in many circumstances, there is a Windows client making it a little more useful than GRE or *vtun*. *OpenVPN* can be found at <http://openvpn.sf.net/>.

Click Click Admin with Webmin

More often than not, administrating a Linux device remotely using a shell can become quite cumbersome and particularly difficult for a novice user to get a whole lot done. Having an easy to use point and click interface available to keep systems ticking over is very useful and there is a great web-based admin tool called *Webmin* which is distributed with nearly all distributions and is certainly available for all of the popular ones.

Webmin provides a web front-end to nearly every service and administrative process on a Linux box, from creating and modifying *cron* jobs, maintaining a MySQL or *Apache* server, adding users or adding domains to our mail server. As each distribution usually installs all of these things in various different places, *Webmin* usually has to be customized so it knows where exactly to look for the configuration files as well as the binaries.

If there is a process that can't be done via *Webmin*, chances are it's something that is done so infrequently it's not worth writing a front-end for, or it's sufficiently obscure that no one who uses *Webmin* actually cares about it. There are literally hundreds of third-party *Webmin* modules at <http://webmin.thirdpartymodules.com/>, so even if something is not maintainable via the



standard *Webmin* installation, there may very well be a module, which can be included to work around that limitation.

As *Webmin* doesn't completely take over the system, it allows us to administrate part of the box over the web, and do the rest using shells. As long as we don't confuse things by trying to modify the same process using two different methods, *Webmin* is a very reliable tool for performing maintenance on device.

Not that the device has to be a million miles away for *Webmin* to be useful, as we can use it to maintain boxes on our LAN through the graphical front-end. Using *Webmin* it's actually very difficult to break a box, and as we're sure many administrators have done, when you're shelled in as root, there's always a day when a space ends up between / and *, or something gets typoed and we overwrite the / filesystem with the contents of /dev/kmem.

Webmin is also very useful when more than one person is require to maintain a system. Not only does it provide a consistent maintenance process, but it also avoids throwing around the root password for the box. While it does not avoid problems where administrators step on each other's toes, it does eliminate many issues where a configuration file may be written in an obscure fashion, or not make it entirely clear as to what a particular option on a system does.

Summary

In this article we've looked at a wide range of ways to maintain a Linux system remotely and hopefully given confidence in particular tools and processes. This is by no means the end of the list of remote administration tools, as there are numerous other less popular ways to maintain systems or reconfigure processes and services. Depending upon one's abilities with the Unix shell, there are some great alternatives to typing then praying, as many have done in the past when hoping that restarting a particular service isn't going to block them out of their server.

As a final note, it's always a wonderful idea to have multiple entry points into a remote system. Never simply rely on SSH, as more than likely, at some point it will die or not restart right, making it mighty difficult to fix without either super-powers or a day out of our life. Having a serial console from somewhere is always an advantage, as is running telnet way up on a high port that no standard port scans are going to find. **LXF**

Webmin offers easy web-based configuration for nearly all Linux network services, including Apache.

NEXT MONTH

With the prevalence of spam ever-growing, FTP (File Transfer Protocol) and filesharing is becoming ever-more important...

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...



Hans Huberland is Rackspace Managed Hosting's Linux expert. Send any Linux system admin questions to sysadminqa@rackspace.co.uk



Firebird

Q Encouraged by LXF45's *Mozilla* article, I downloaded the Windows & Linux versions of MozillaFirebird-0.6.1.

The Windows version was installed on Win 98/2 by clicking on it, so I had it working in about five minutes. The only problem was that a shortcut/icon could not be installed.

The Linux version was extracted by `tar xvzf` as usual and the *MozillaFirebird* directory entered. In both SUSE & Red Hat I got no further, as there was no config file and the arcane instructions were not understandable. Some hours later, after various random fiddling, the Red Hat version worked via a further *MozillaFirebird* instruction within the *MozillaFirebird* directory. How should this program be installed?

I had not noticed you had this on the disc. I have now managed to load this OK and get it working.

Two oddities:

1 I can generally only get these progs to install by copying them to my home directory and then installing from there. Following the instructions on LXF p107 I find never works. Perhaps operator error?

2 `./configure` still failed but I found the prog seemed to install itself within the original folder without further work after tarring.

Joe Lamb, via email

A Once you untar your *MozillaFirebird* tarball, you can run *Mozilla* from within there using the

`./MozillaFirebird` command. Copying the entire directory to `/usr/local/mozillafirebird`, then symlinking the scripts to run mozilla into `/usr/local/bin/` will allow you to run *Mozilla* with simply the `mozilla` command. The *Mozilla* tarballs are distributed as binaries, rather than something that you will need to



The Mozilla project distributes binary tarballs of *Mozilla* and *Mozilla Firebird*, but they don't install like standard packages.

configure and compile yourself. mozilla.org has documentation on installing the tarballs it distributes on a system, although depending on your distribution you may manually have to add menu options for KDE or GNOME.

Learning & Linux

Q I work for a small community school in the US. Recently my administrator purchased some systems from a company called New Internet Computer (against my advice, but they were very inexpensive).

The systems have no hard drive about 64MB RAM, Ethernet, modem, video, USB and sound. It requires a boot CD to load the OS and connect to the Internet. I used the bootable demo CD that came with one of your issues of *Linux Format* and it worked. However it was very slow because of all of the options that ran on it. I need to know how (or whether it is possible) to customise a bootable CD that

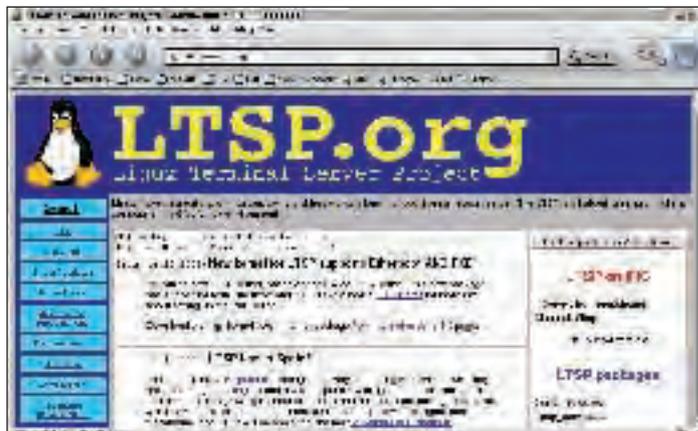
runs the bare minimum modules plus a browser for Internet access.

Next, I would like to know how to have it connect to a Linux server to grab and execute basic software like word-processing and/or spreadsheet and store files to that server so that students can get more use out of it more than just Internet access. I hope you can help me.

Tommy Bordain, via email

A These systems are more comparable to terminal servers, rather than any sort of real workstation. Having no local storage will make life rather difficult, as you'll either have to run everything off a CD or over a NFS mount. A great bootable CD is Knoppix, which comes with all sorts of things from word processors and web browsers, to development software.

A great project to make use of these devices would be the Linux Terminal Server Project, at <http://ltsp.org/>. With a large central file and application server, which all of



Linux Terminal Server Project is a great way to make use of old or under specced systems, as we can build a large backend server to run apps on.

the thin clients boot off over the network. You can run them as loaded, or as unloaded, as you like, offering X applications over the network. LTSP offer a great distribution to build the server, so it's very easy to get the system up and running. All that is really required is to obtain ROMs for the NICs to permit them to boot off the network if they do not already.

On the plus side, you can buy hard drives exceptionally cheaply now, so any savings could be invested in buying real storage devices for these boxes. Of course, depending how

poorly they are specified on the CPU front, it may not be worth the effort.

Glub about glub?

Q I've been reading your magazine for the past six months. I like it a lot. I installed *glib-2.2.3*. Then I tried installing ATK so that I could install GTK. When I used the configure command for ATK, I got the following message:

checking for GLIB - version >= 2.0.0... 'pkg-config --modversion glib-2.0'

returned 2.2.3, but GLIB (2.0.6) was found! If pkg-config was correct, then it is best to remove the old version of GLib. You may also be able to fix the error by modifying your LD_LIBRARY_PATH environment variable, or by editing /etc/ld.so.conf. Make sure you have run ldconfig if that is required on your system. If pkg-config was wrong, set the environment variable PKG_CONFIG_PATH to point to the correct configuration files no configure: error: GLIB 2.0.0 or better is required. The latest version of GLIB is always available from. If GLIB is installed but not in the same location as pkg-config add the location of the file glib-2.0.pc to the environment variable PKG_CONFIG_PATH.

I modified both the LD_LIBRARY_PATH and the PKG_CONFIG_PATH to include the directory that contained the .pc files (/usr/local/lib/pkgconfig). I used the export command to do this and when I echoed the variables, the only directory listed was /usr/local/lib/pkgconfig. I also edited the ld.so.conf. file to include this directory and ran ldconfig. I still get the same error while trying to

configure ATK. I can not find glib-2.0.6 so that I can remove it. I am not even sure if the .pc files are what is being looked for or if there are configuration files somewhere else. I have seen similar posts on Linux help boards for what appears to be the same problem, but I have found no good answers.

AJ Monatesti, via email

A You most likely have RPMs or other packages for *glib* installed already, which are going to cause conflicts. You didn't tell us which distribution you were running, but generally you don't want to mix sources and packages on the same system unless you have plenty of experience maintaining libraries and headers on a Linux box. Removing the *glib* packages will require that GTK+ be uninstalled too, which in turn will want to take the majority of GNOME and other applications such as *The GIMP* with it.

You probably don't want to try to install *glib* and GTK+ unless it's distributed as an update or upgrade for your particular distribution, as it's going to cause lots of other problems. Moving from *glib* 2.0 to *glib* 2.2 will require you to recompile all of GNOME and any other binary that is linked to *glib* or *gtk+*, which is quite an undertaking. Unless you've got a pressing reason to upgrade *glib*, you might just want to give it a miss.

Respect my authoritah!

Q I have a Red Hat 9 setup with just a couple of users... X starts when the machine boots and launches GNOME. The other day I logged in as me and got an alert saying:

You have logged out within 10 seconds. If you didn't do this on purpose you may be out of disk space. See ~/.xsessionerrors

I logged in as root and saw that there were now 2GB of music files in my home directory – I guess I need to use *gnutella* a bit less! So I moved most of them out of there and tried to log in as me again. Same alert. So I checked *.xsession-errors* (which I should have done first) and which had a line about not being able to find/open the file *~/.ICEAuthority*. That file is there but it's corrupt.

A QUICK REFERENCE TO: Links

It's often very useful to be able to quickly access files or programs without having to wade through layers of directories to find what you're looking for. Rather than having to remember your way through a mess of files and sub-directories, it's much easier to create links from the original file to somewhere a little more convenient.

A link is simply a special type of file that refers to another location on the system, although links can either link to another filename, or to a specific place on a filesystem, known as an *inode*. The former, known as symbolic links, like one file to another by use of its filename. If we do:

```
In -s docs/howto/links.txt links.txt
```

whenever *links.txt* is accessed, the VFS layer will really access *docs/howto/links.txt*. Symbolic links can be created absolute or

relative to their source, so if we move the 'links.txt' symbolic link into a directory known as 'stuff', it will then try to look for 'stuff/docs/howto/links.txt'. To counter this, we can create an absolute link:

```
In -s ~/docs/howto/links.txt links.txt
```

Now, wherever the symbolic link is moved to, it will always refer to the original file correctly. If we rename, or delete, the original *links.txt* file, the link will no longer work, and is known as a 'broken link'. When we try to access a broken link, the VFS will return a 'file not found' error, just as if the symbolic link wasn't there at all.

Within a filesystem, we can also create a hard link. Rather than referring to the filename, which may change, hard links refer to the location of the original file on the filesystem. Each filesystem has a finite number of locations from

which a file may start – those inodes as mentioned before – and a hard link simply points to this place. A hard link is effectively an alternative name for a file on a filesystem, since we can do whatever we like with the original location.

```
In ~/docs/howto/links.txt links.txt
```

We can delete and rename the file, and the hard link will continue to function properly. When a file has been hard linked there is nothing to differentiate between the original filename and the link, so both must be deleted before the space used by the file data is reused by the filesystem. However, the caveat of hard links is that *they may only exist on the same filesystem as the original file*, so if /home and / are separate partitions, then it follows that we can't create a hard link from */home/david/links.txt* to */links.txt*.

FREQUENTLY ASKED QUESTIONS FIREWALLS

FAQ HOW TO I STOP REMOTE MACHINES CONNECTING TO MY COMPUTER?

There are two different ways to block access by remote machines to your computer. The first is to restrict access within particular services, so that they deny access to specific hosts. Many use `/etc/hosts.allow` and `/etc/hosts.deny`, and you might set yours up as;

```
hosts.allow;
```

```
ALL: LOCAL
ALL: 192.168.
```

```
hosts.allow;
```

```
ALL: ALL
```

This will block access to all machines, except those with an IP starting **192.168**, or if the connection is local to the machine. However, many services do not take advantage of `hosts.allow` and `hosts.deny`, so we need to use something else.

Depending upon your Linux kernel version, you can either use `ipchains` or `iptables`. Linux 2.4 users should use `iptables`, as it has far more advanced firewalling capabilities than `ipchains`.

We can block access to our machine on a specific interface, such as `ppp0`, using the following:

```
# iptables -A INPUT -i ppp0 -j DROP
```

However, this will drop all connections, including those which did not originate externally, so we will need to allow connections which we started to work:

```
# iptables -I INPUT -i ppp0 -m state
--state RELATED,ESTABLISHED -j ALLOW
```

We can check our firewall rules with:

```
# iptables -nvL
```

`iptables` requires kernel modules to be compiled, and a user-space application 'iptables' to be compiled. The latest release of `iptables` can be found at www.netfilter.org.

FAQ CAN I PERMIT ACCESS TO SPECIFIC PORTS?

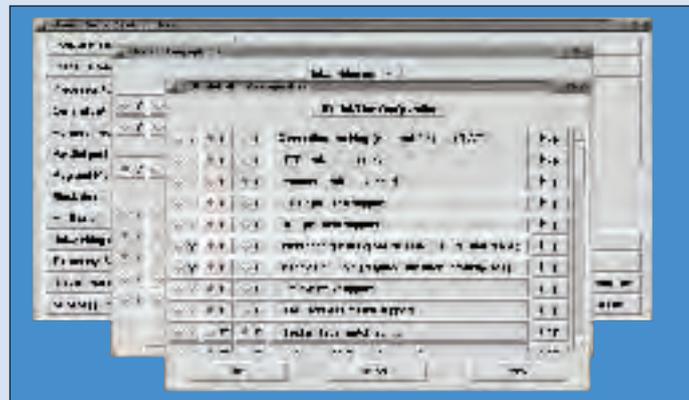
Using `iptables`, we can allow packets belonging to a specific IP protocol through, such as UDP or TCP. It's important that we allow packets before they are dropped, so we have to use `-I` to insert it before our DROP rule thus:

```
# iptables -I INPUT -p tcp --dport
22 -j ACCEPT
```

which would allow access to the SSH port, **22**, from anywhere.

FAQ CAN I SHARE MY INTERNET CONNECTION WITH OTHER SYSTEMS ON MY NETWORK?

In order to share the network connection, we need to use IP



The Linux kernel has numerous firewalling options built right in, making the construction of a secure Linux-based firewall very painless.

Masquerading. Assuming the network is set up correctly, and the IP of our machine is the gateway for the other machines connected to the network, then we only have to enable IP forwarding, then setup masquerading. To enable IP forwarding, we just need to do;

```
# echo 1 >
/proc/sys/net/ipv4/ip_forward
```

And for masquerading, assuming `ppp0` is the interface connected to the Internet:

```
# iptables -t nat -A POSTROUTING
-o ppp0 -j MASQUERADE
```

FAQ HOW CAN I REDIRECT A PORT TO A BOX ON MY NETWORK?

In order to redirect a port, we have to use the DNAT, or destination NAT, capability of `iptables`. This rewrites a packet so that the destination IP of the packet changes from the IP of our external interface IP, to an IP of an internal machine, so the packet is routed to it.

If we wanted port 22 on the gateway connected to the Internet to be redirected to a machine on 10.1.1.2, we could do:

```
# iptables -t nat -A PREROUTING -i
ppp0 -p tcp --dport 22 -j DNAT --
to 10.1.1.2:22
```

It is perfectly possible to redirect one port onto another; so rather than being redirect to port 22, we could redirect it to 10.1.1.2:3452 instead if we so desired.

◀ If I try to list it I get:

```
ls: /home/jim/.ICEauthority:
Input/output error
and another file is also corrupt:
/.DCOPserver_machines.address.co
m_0.
```

Can't stat them at all.

I read a bit about `.ICEauthority`, what it's for and that it can lock if

Posting to the forum The LXF online community

Not only do our popular forums at www.linuxformat.co.uk have sections dedicated to your technical queries, hardware, programming languages and general help; but also there's always a lively discussion going on!

you crash, but got no information on restoring it. I tried `iceauth` and can do:

```
iceauth -f /home/jim/.ICEAuthority
iceauth: creating new authority file
/home/jim/.ICEAuthority Using
authority file /home/jim/.ICEAuthority
iceauth>
```

but then what? If I save this 'new' `.ICEauthority` file that `iceauth` says it created with the `exit` command, the corrupt one is still there. I have copied my data and prefs, so if you can tell me how to get rid of corrupt files like that, I could just remove my old account then make another – but restoring this file would be better!

Any pointers gratefully received!
Jim Burton, via email

A If you received an I/O error from doing an `ls`, then it suggests there is a bigger problem. If you ran out of disk space, you may simply want to reboot the box into single user mode and force a `fsck` of the filesystem to ensure that everything is consistent. You'll also want to run `dmesg` after you see the I/O error to find out if there is a filesystem issue which is being logged by the kernel.

You can safely delete `.ICEauthority` from your home directory, assuming it will let you, as it will be recreated if required. Of course, if you don't delete it before trying to do anything, you're going to see the same error, as you will be hitting the same inode with each filesystem request. If you can't delete it, a `fsck` should solve that problem.

Printer pleas

Q I've got a big order in and this problem is getting in the way. I'm currently running SUSE 8.0 Pro on a PIII connected to a couple of printers: the main donkey being HP8000DN, connected at the moment through a parallel port, running CUPS as the printer spooler/daemon/whatever.

The snag is: often on a large PostScript file the print job stops, the small screen on the printer flashes 'Ready', then flashes 'Processing Data', then 'Ready...' in a loop. It doesn't matter how long you wait, nothing is printed. Doing `lpq -Phd` brings up the current status of the printing which is still active. `lprm [job-number]` clears

FAQ WHAT'S THE DIFFERENCE BETWEEN SNAT AND MASQUERADE?

If we have a connection on a dynamic IP, then we need to use the MASQUERADE rule, as when the connection goes down, the connection tracking system will automatically flush all of the rules belonging to the interface as a security measure.

However, if we have a static IP, this will stop TCP connections from resuming, which is rather annoying. By using a SNAT rule, we can have all the features of the MASQUERADE rule, but if the connection goes down, then all tracked connections will remain in the list.

FAQ CAN I LOG UNAUTHORISED CONNECTIONS TO MY COMPUTER?

You can log packets through *syslogd* using the LOG rule in *iptables*. Usually, two separate rules are used for each set of packets to be dropped, one for the logging, and the other to drop the packets. To log all incoming connections, then drop them, we would do:

```
# iptables -A INPUT -i ppp0 -m state --state NEW -m limit --limit 1/sec -j LOG
# iptables -A INPUT -i ppp0 -m state --state NEW -j DROP
```

By specifying a log limit, we can avoid attacks which attempt to fill up our /var partition with dropped packet logs.

the printer, the printer is shown as being idle and able to print, but it is just not up for it. Send the printer any job you like and nothing gets done. Turn the printer on and off, it just shows 'Ready', but does not print. *lpq* will show a list of jobs, the first one being active.

At this stage, I settle for rebooting, which can't be the solution. Strangely, if I leave the printer on, the damn thing starts to print the job, the whole job that it fell down on in the first place as either it boots down or up! So, something must be in the spooler somewhere, still getting pushed to the printer.

So, without rebooting just for a duff print job, what's the solution?

Flush the printer cache, CUPS, unload and load CUPS? I've tried killing *cupsd* and then restarting but the printer just does not want to know. Any advice...?

Daithidh MacEochaidh, via email

A You may want to experiment with restarting CUPS from the command line with:

```
# /etc/init.d/cups stop
# /etc/init.d/cups start
```

Of course, why the print job stops in the first place is another issue. Does the printer have enough buffer to handle the amount of data you're throwing at it? Restarting the printer will generally not help, so you'll need to look at the Linux end of it all. Does the printer manual indicate what the problem is when you see the loop of the status lights? It could indicate a misconfiguration at the CUPS side of the process.

Cordless confusion

Q I recently bought a nice Gyration Ultra cordless suite (wireless mouse and keyboard) to plug into my dual boot Mandrake 9.1 / Windows XP desktop. I didn't anticipate any problems, as Mandrake has served me well with hardware recognition until now, and any hardware that doesn't even supply a Windows driver disk must be pretty sure to work just about anywhere. And it does, except for the five or so seconds I need it to work when I'm picking my OS at boot!

The keyboard works fine in the BIOS, and in whichever OS I eventually boot in to, but for some reason, when the *Grub* appears, the wireless receiver cuts out, and I'm left powerless to pick where I want to boot. The set basically consists of a USB wireless receiver, a keyboard and a mouse with some clever gyroscopes inside. I'm not particularly concerned with the fact that the mouse doesn't work with Mandrake, as I'll try out 9.2 and see where that takes me (Red Hat 8 and JAMD 0.6 seem to have no problems), it's just that I need the keyboard to work with my bootloader. Even more bizarrely, when I plug in the receiver to my Dell laptop (running Mandrake 9.1) the keyboard *does* work in the bootloader. Any ideas? Thanks for a great mag!

Jim Rutherford, via email

A Linux distributions will either use *Grub* or *Lilo* as a boot loader. Unfortunately, the former is known to not work happily with USB keyboards, so the only real solution is to migrate your Mandrake 9.1 system to the *Lilo* boot loader to allow it to work with the keyboard. You may want to compare the boot loader configuration on your laptop to that on your desktop system, so that you can check out any configuration differences there might be.

Sadly, as *Grub* is not really a part of the main Linux installation, it follows that the Mandrake hardware detection code will not play any part in how *Grub* works.

Many mailing lists have documented examples where *Grub* fails to work with USB keyboards and *Lilo* will. Until someone fixes *Grub*, then you'll sadly be stuck with *Lilo* if you want to continue with your wireless USB keyboard.

Building KDE

Q I would like to build KDE 3.1.3 but don't seem to be able to get past the arts-1.1.3 configure script which exits with the following error message:

checking for Qt...

configure: error: Qt (>= Qt 3.1.0) (library qt-mt) not found. Please check your installation! For more details about this problem, look at the end of config.log. Make sure that you have compiled Qt with thread support!

As far as I can tell the *qt-mt* library is available:

```
root@starbug:/home/steve > ldconfig
-v | grep libq
libqui.so.1 -> libqui.so.1.0.0
libqt-mt.so.3 -> libqt-mt.so.3.1.2
libqtmcp.so.1 -> libqtmcp.so.1.0.0
```

Qt was configured with the following options:

```
./configure --system-zlib --qt-gif
--system-libpng --system-libjpeg
--plugin-imgfmt-mng --thread
--no-stl --no-xinerama --no-g++
--exceptions.
```

I tried upgrading my compiler from *gcc-2.95.3* to *gcc-3.3* and am currently using *gcc-3.3.1*. I am able to build *qt* successfully with any of these compilers, but in this case I am not able to get past the *arts* configure script. I have tried both *qt-x11-free-3.1.2* and *qt-x11-free-*

3.2.1, but each time *arts* fails at the *./configure* stage.

I have compiled plenty of other packages using *gcc-3.3.1* (including kernel-2.4.22) and whilst 3.3.1 is a good deal more fussy than 2.95.3, I don't think the compiler is what's causing the problem.

I built KDE 3.0, 3.1 and 3.1.1 (which I am still using) without any trouble, but KDE 3.1.2 and 3.1.3 have got me beat. If you guys at *Linux Format* are able to help, I will be mighty grateful!

Steve Roper, via email

A To compile something against a library, you require both the library itself, as well as its headers. Even if you have the library installed, the tests for it will fail, as it compiles a test program which requires the headers from the *-devel* package. If you obtain the *qt-devel* packages from the same place you got the *Qt* library package, you will be able to build KDE properly.

As you've got *Qt* installed using packages, you may want to search for KDE packages for your distribution to avoid mixing source code and packages on one system.

Submission advice

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

- Please be sure to include any relevant details of your system. 'I can't get X to work' doesn't really mean anything to us if we don't know things like what version of X you are trying to run, what hardware you are running on.
- Be specific about your problem. Things like 'it doesn't work' or 'I get an error' aren't all that helpful. In what way does something not work? What were you expecting to happen? What does the error message actually say?

- Please remember that the people who write this magazine are NOT the authors or developers of Linux, any particular package or distro. Sometimes the people responsible for software have more information available on websites etc. Try reading the documentation!

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

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

ANSWERS



Deleting files

Q I have a directory with a large amount of files that I would like removed.

However when I try to remove the files using `rm *` inside the directory I get the following error:

```
/bin/rm: Argument list too long
```

I know I can delete this entire directory directly with a `rm -fr` but I would like to delete some of the files using a wildcard without effecting the base directory. How can I go about this?

Samuel, via email

A Because there are so many files in this directory, when you try to do a `rm *`, you try to pass every single file as an argument to the `rm` command, more arguments than it can handle. There are a few possible solutions to this problem but I will go with the one I feel is most useful for any program which only takes a limited number of arguments. To get around this problem there is a very useful utility called `xargs`. `xargs` builds and executes command lines from standard input.

Using the following command we can create a list of all the files and pipe it to `xargs rm` which will delete these files one by one:

```
#ls -l | xargs rm
```

If you would like to only delete specific files indicated by a wildcard – for example all `.txt` files – you can do the following:

```
#ls -l *.txt | xargs rm
```

Router/firewall

Q I have a Linux server running Gentoo (the Linux distribution, as opposed to the file manager on this month's LXF coverdisc) on a network and I need to share the server's Internet connection with my other Linux boxes on the network.

My server is both connected to the Internet and the network at the same time. How do I make my sever act as a router/firewall?

David, via email

A You've given me very little information to go on, but I assume that you will be using non-routeable IP addresses on your private network. If so, then iptables will be able to do this for you.

To enable Network Address Translation (NAT) or masquerading, you will need to issue the following command:

```
iptables -A POSTROUTING -t nat -o ppp0 -j MASQUERADE
```

However, you will also need to ensure that your kernel will allow routing between interfaces, this is done setting `/proc/sys/net/ipv4/ip_forward` to 1, either by typing:

```
echo 1 > /proc/sys/net/ipv4/ip_forward
```

or by editing `/etc/sysctl.conf` and looking for the value `net.ipv4.ip_forward` and again setting it to one. The latter makes the change permanent.

Please, please, please remember that this *does not* provide any protection to your network. I would highly recommend that you add some more rules to prevent people from spoofing their way into your network. The following two lines of code will allow all outbound connections but will only allow packets from already established connections to come into the network:

```
iptables -A FORWARD -i ppp0 -o eth0 -m state --state ESTABLISHED,RELATED -j ACCEPT
iptables -A FORWARD -i eth0 -o ppp0 -j ACCEPT
```

Ghastly Gallery

Q I've recently been struggling to install *Gallery* on my Mandrake server. So far, so good, but when I try to add photos I get the following error:

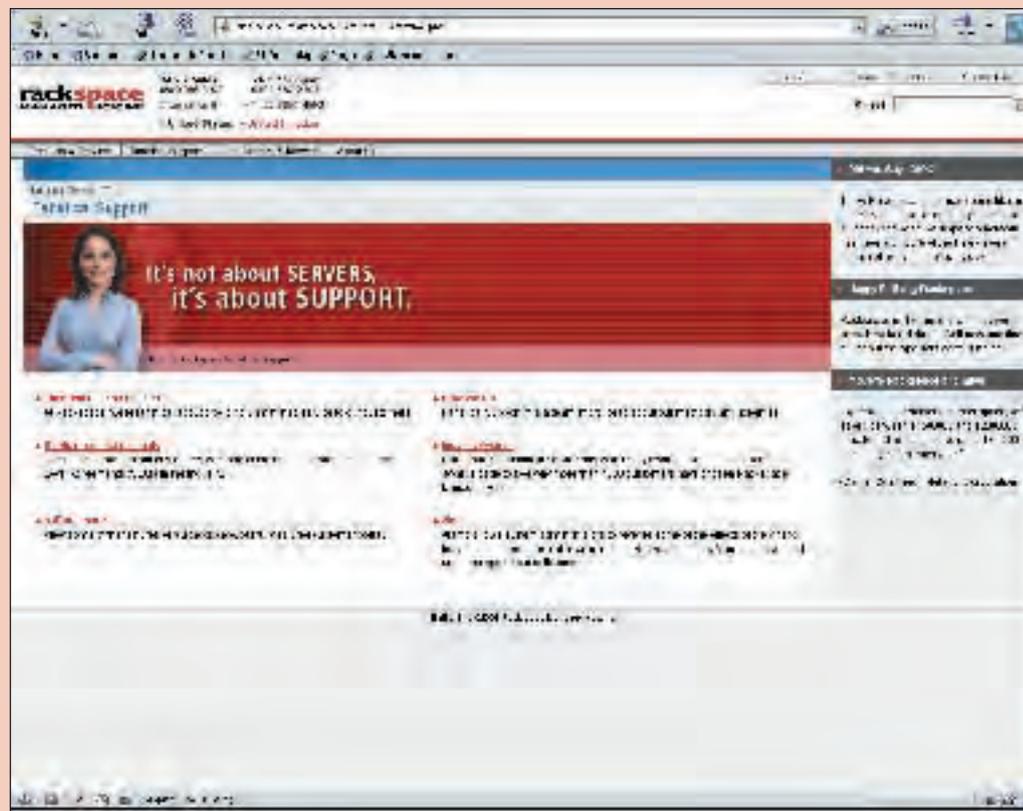
Warning: open_basedir restriction in effect. File is in wrong directory in /home/httpd/vhosts/www.domain.com/httpdocs/gallery/platform/fs_unix.php on line 53

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 and see this month's *Linux Pro*.



WIN A SHARP ZAURUS PDA FROM RACKSPACE MANAGED HOSTING

www.rackspace.co.uk

Every month, the best question related to Systems Administration that a LXF reader sends in wins a Sharp Zaurus SL-5500 running Linux.

With industry support from the likes of IBM, SAP, Oracle, and Sybase, the Zaurus features a QWERTY keyboard behind a sliding cover. The 65,536 colour screen is 240x320 resolution. An Intel® StrongARM® 206MHz processor and 64MB memory makes the Zaurus ideal for downloading software or playing MP3 and MPEG1. The docking station that fits into its USB connector, allowing users to synch with their desktop, schedules, address

book and ToDo lists. The email app supports POP3 or IMAP email servers. Write and send emails and browse the Internet with a compatible modem card, or using the infra-red port and a mobile phone; or any replied, forwarded, or composed emails are transferred to your outbox the next time you synch. SD and CF slots allow upgrades and addition of software.

Email all your sysadmin questions and problems to Hans Huberland: sysadminqa@rackspace.co.uk Every published query on these system admin pages wins a Rackspace T-shirt.

It's only 74mm wide by 138mm long and a svelte 18mm thick – so the Sharp Zaurus really is Linux in your pocket!



★ Star Question – PDA winner!

This issue's lucky winner is **Max** – your new Zaurus PDA will be with you shortly!

Shifting IP address

Q I'm being forced (by my choice to change ISP) to change the IP address of my server. This server hosts about 300 domains and I've been warned by my potential new ISP that we will have to wait for DNS propagation on all those domains once we change to IP address. I don't really understand why this process needs to take 24 hours. Surely if I change the DNS settings, then anybody looking to my DNS server for information will get the correct information? I really don't want to have to tell my clients that their sites will be down for a whole 24 hours.

Max, via email

A The problem is a fairly common one when changing IP address. Fortunately, by the sound of your email you have not made the move yet, so there is still time. Further, because you run the DNS server yourself you already have a lot more control than users who are relying on their ISPs.

Every DNS zone has a time-to-live (TTL) value associated to it. This tells DNS servers that pick up this zone information from you how long it is valid for. By default this is 24 hours (86400 seconds). This is done to reduce Internet traffic and to reduce the workload on DNS servers of popular domain names. The DNS server will then give these results out to any requesting client until the TTL expires, at this time it will need to look up the information from your DNS server again.

That said, you can change the TTL on all your domains to a low value (such as 5 minutes) and then wait 24 hours for all the other DNS servers' caches to expire. When a client request the IP for one of domain names, it will get the same info but with a 5 minute TTL. You can now change the IP address and put the TTL back to 24 hours. This means that all the servers with the old IP will only keep it in cache for a maximum of 5 minutes after you make the change. The one caveat is that some ISPs choose to override this TTL with their own setting regardless and there is nothing that can be done about this – if the site in an intranet site or similar and appeals only to a select few on a particular network, you may want to check with their ISP if this is the case.

This is the code that appears on line 53:

```
function fs_fopen($filename,
    $mode, $use_include_path=0) {
    return fopen($filename, $mode,
    $use_include_path);
}
```

Any ideas? Can you help?

Karl, via email

A You should make the changes in your main php configuration file. This file should be /etc/php.ini. Searching through this file should reveal an open_basedir value. This value limits the maximum file size that can be operated on the server (including uploads). This is a global setting. If you wish to set this locally then you need to add a line containing:

```
php_admin_value open_basedir
"some_path"
```

to your virtual host section of your apache configuration. However, the ability to set this per host appears to be broken in PHP 4.3 so you'll need to use the global setting if you're using this version. For a lot more useful information visit <http://bugs.php.net/bug.php?id=19292>

Grokking the Ghost

Q I'm trying to create a backup disk of my current Red Hat Enterprise Server setup.

I'm using Norton Ghost to create the image. It will ghost fine but after ghosting Lilo displays "LI" only

and the machine seems to be hanging. The problem is fixable by booting to the Linux floppy or Rescue environment and rerunning lilo which rewrites the Master Boot Record. Why won't the ghost with Lilo work the first time round? Thanks for your suggestions!

Josh, via email

A Lilo's LI error implies that the first-stage loader was able to load the second-stage loader but was not able to execute it. It's normally caused by a drive geometry change or by boot.b being moved. As you are ghosting to a new drive my guess would be that the geometry of the two drives may not be exactly the

same. As the second stage boot loader falls outside the first 512 bytes of the drive (The MBR), this could be in an entirely different physical location on the new drive. As lilo is still pointing to the old location, it cannot execute the second stage boot loader.

The alternative is to use a boot loader like Grub, which does not require any reload after a ghost as its second-stage boot loader is not fixed by drive geometry. You can find more information on Grub at www.gnu.org/software/grub/. Bear in mind that as you are running Red Hat Enterprise Server, loading Grub may invalidate your support. Check your terms and conditions!

« Installing

Q I have a Suse 8.2 running on an AMD athlon desktop PC. I am trying to install *MPlayer* from the *LXF* Disk. I use `./configure --enable-gui && make && make install` and it compiles fine. Then I can open AVI files from *MPlayer* (command line). But I would like to use the GUI (*gmplayer*). When I try to do this I get a message "gui.conf missing". I have read the documentation about how to configure, but it doesn't speak English!

I have tried the other RPMs, but they simply do not work, stating that *MPlayer0.90* is a dependency.

Please could you shed some light on this, I am only a newbie!

Chris Taylor, via email

A If you have RPMs for *MPlayer*, then you may want to do a `make uninstall` on your source tree, and then install all of the packages. You should be able to do `touch ~/.mplayer/gui.conf` and then run the GUI version of *MPlayer*. As a self-proclaimed newbie, sticking to packages is probably a good option to start with, as building from source code can often introduce more problems than you really want to deal with until you're a bit more skilled.

Throw off your mental chains...

Q I have just bought a *Netgear DG814 ADSL Modem/Switch* – a nice piece of kit which means I can now share my broadband connection with my Linux machines and not be dependent on a USB connection to the the Windows box.

I configured the modem to be a DHCP server for my small network and it works fine. I presume DHCP has to be the case because the DNS server addresses form BT are dynamic. `/etc/resolv.conf` is dated stamped at boot time.

The trouble is that because of DHCP, my old hosts file is now wrong and it would be nice to be able to rcp, mount NFS etc without having to check the IP address first. To be fair, I think the IP allocation is fairly static and for a network of 4 we are not talking a major problem. But in principle I should be able to either update my hosts file or use



MPlayer is a wonderful skinnable media player, but some users have found that it can be quite a pain to install.

some kind of local DNS to make it work properly. The *DG814* has a screen called 'Attached Devices' with all the info required IP address, Device Name – even Mac address – but I don't know how to get to it. (I have asked Netgear support; and will let you know if I hear anything useful). If you have any suggestions, they would be greatly appreciated.

Howard Jones, via email

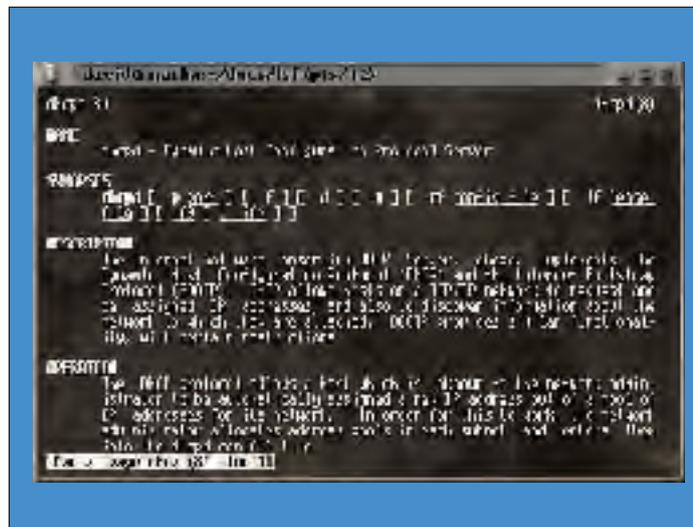
A If you only have a few devices, then why not avoid using DHCP at all and simply use static IP addressing for everything? There isn't really a simple or standard way to obtain the DHCP information from that device, although with many DHCP servers, you can either statically assign IPs based upon the MAC address, or have the DHCP server

push updated records out to a DNS server on the network with the appropriate name server information.

For a larger network, running a dedicated DHCP/DNS server is a practical solution, since you may be dealing with hundreds of hosts, if not more. However, for a smaller network, using static IP addresses is considerably less work.

Shell shy

Q Having played around with many Linux distros over the last few years, presently running *Mandrake 9.1*, one aspect of Linux still eludes me: software installation. I read your instructions regarding installing from tarballs, have Linux reference books; maybe I'm a bit slow but I still can't figure out 'how to'.



The DHCP server from isc.org allows us to push out DNS updates to BIND when a new DHCP lease is offered to a client.

Do I use one of the editors to type in command lines, `cd` to directory? Press **Enter** after each line etc? Whatever I do, it doesn't work for me, so I'm obviously doing the wrong thing. I am left staring at the treasure trove of software you provide, unable to try any of it.

Les Mitchell, via email

A *Linux Format* has run some articles in the past on basic shell usage and how to maintain a Linux system. Past issues can be ordered on page 97 and found online at www.linuxformat.co.uk/. There are also various other sites aimed at Linux Newbies, including www.tuxfiles.org/ and www.linuxnewbies.org/, both of which have great information for those starting off with Linux.

New hard drive

Q I wonder if you can help me start with Linux. I have a *Hewlett Packard Pavilion 6260*, about three years old with *Windows 98* loaded from a retailer's disc. It has an AMD K6 processor performing 500 megabits per second and 64MB RAM. I am hoping to replace the *Maxtor 10GB* hard drive with a similar modern substitute and then load the Linux onto the new hard drive. Is this feasible? If not, what extra equipment will I also need? I am also intending to upgrade the RAM to 128MB.

Anthony Stephens, Oldbury

A You will be able to install both drives in your system at the same time and partition the new device and migrate all of your data over to it. Moving Linux from one drive to another can be somewhat difficult, as you need to boot off a rescue disk or boot it into single user mode to copy the root filesystem. You may simply want to install a completely fresh Linux system onto the new drive and then copy the data over from the old hard drive. Just remember that it's essential to back your data up to CD or tape, just in case you repartition the wrong device and lose everything.

On the hardware front, you may just want to get a load of new RAM for the box. 512MB of PC133 is exceptionally cheap now and upgrading to 128MB probably isn't going to give you a whole lot of extra performance. LXF

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 program 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).



December 2003

Product code:
LXFB0047(cd)
LXFDB0047(dvd)

DVD HIGHLIGHTS:
OpenOffice.org (PPC),
MailScanner, XFCE,
LDP, PuppyLinux,
Slurm, fphoto,
Battlemech, XMMS,
Fandango, CVSToys,
Scrubber, Sylpheed

MAGAZINE FEATURING:
Mandrake 9.2, Linux Expo UK
report, Lazarus, CVS, Xinit
SharqStation, MTStudio, SUSE
9.0 review, Scanners tutorial,
Newsreaders roundup, IBM's
Linux strategy, Nessus

CDs HIGHLIGHTS:
Mandrake 9.2,
OpenOffice.org,
GNOME 2.4



November 2003

Product code:
LXFB0046(cd)
LXFDB0046(dvd)

DVD HIGHLIGHTS:
RUNT, RPM Analyser,
Linux Easy Access
Keyboard, VCool

MAGAZINE FEATURING:
The new 2.6 kernel, Java IDEs
Roundup, PHP 5 preview,
Linux keyboards, Aethera, The
Grid, OOo 1.1, Armari Pro3D,
Regular Expressions tutorial,
New-look SUSE, ZoneMinder

CDs HIGHLIGHTS:
Gentoo Linux, Seapine
Surround SCM FULL
VERSION, Nano, Mass
Rotate, BZFlag,
Jahshaka, Kroupware,
ZoneMinder, BatchGLIMP,
ConfluxGroupware



October 2003

Product code:
LXFB0045(cd)
LXFDB0045(dvd)

DVD HIGHLIGHTS:
Topologilinux, JRE SDK,
Wifiscanner, Knokisync,
Squirrel SQL Client

MAGAZINE FEATURING:
Speed up Linux, Red Hat
Package Management part 2,
Server School tutorial,
OpenGroupware.org review,
Which webserver, What On
Earth is BSD? Stop spam

CDs HIGHLIGHTS:
FreeBSD4.8, K3b,
Kognition, Beneath A
Steel Sky, ScummVM,
Thunderbird & Firebird,
Java2, Java Runtime
Environment, Scribus
1.0, LM Sensors



September 2003

Product code:
LXFB0044(cd)
LXFDB0044(dvd)

DVD HIGHLIGHTS:
JAMD Linux, LimeWire,
MoviX, KOffice 1.3,
Scribus, LBreakout2,
Mozilla 1.4, Opera 711

MAGAZINE FEATURING:
40 pages of tutorials, SuSE
Enterprise Desktop 1 review,
Text Editors Roundup, Linux
Sound feature, Seapine SCM,
JBuilder9, Navaho TeamCAT,
Newware Thin Client review

CDs HIGHLIGHTS:
MULTIMEDIA SPECIAL
GNUSound, TuxPaint,
Cinelerra, Realsoft 3D,
AlsaPlayer, TreePad,
Ecamegapedal, Grip,
Kino, Xine, TiMidity++



August 2003

Product code:
LXFB0043(cd)
LXFDB0043(dvd)

DVD HIGHLIGHTS:
GnuCash, 3dDesktop,
Slackware Live CD, LDP,
Astaro Security Linux,
SharpROM v3, Cinelerra,
ScummVM, TrustixSL

MAGAZINE FEATURING:
Networking feature, VMware
review, CD-burner roundup,
Evesham ELinux desktop PC,
Opteron servers, Linux training,
TextMaker 2002, Netvault 7,
GnuCash, 8 beginners' pages

CDs HIGHLIGHTS:
Complete Networking
Toolkit, SpamAssassin,
AGT, GuardDog, Jwall,
Nagios, America's Army,
Samba (stable & dev)



July 2003

Product code:
LXFB0042(cd)
LXFDB0042(dvd)

DVD HIGHLIGHTS:
Opera 7.1 beta, GAIM,
ABCDE, Pingus, KDE-
Sounds/Aqua Fusion,
LDP, AlsaPlayer, KOffice,
RUTE, Seclude, Freevix,
DamnSmallLinux

MAGAZINE FEATURING:
All about X, MySQL review,
Sharp server review, WINE
explained, Steganography,
Distributed Computing, Red
Hat Enterprise AS, Choosing
a distro, GCC 3.3 review

CDs HIGHLIGHTS:
Morphix, GuardDog,
Nvidia XDrivers, GCC
3.3, TORCS, Linux
Cookbook, SETI

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 FORMAT

www.linuxformat.co.uk



SPECIAL CHRISTMAS OFFER!



SAVE 40%* ON THE PERFECT CHRISTMAS GIFT!

A subscription to *Linux Format* makes a fantastic Christmas gift for a friend, relative or even yourself!

- ❄ Save over **£31*** on the price you'd pay in the shops
- ❄ **FREE**** gift card for you to give on Christmas Day
- ❄ **FREE** delivery direct to your door
- ❄ **Never miss an issue!**

LINUX FORMAT BRINGS YOU:

- 1** Essential reading for all Linux users
- 2** 20-page *Linux Pro* supplement with every issue
- 3** Three CDs or a DVD packed with Linux software
- 4** Latest news, latest reviews and the best Linux tutorials around

Taking out a subscription or giving a gift is easy. Simply call **0870 444 8645** or complete the order form opposite and send to the freepost address.

*40% Saving is based on 13 issues bought on the UK newsstand, this offer is only available to delivery addresses within the UK. £31 saving based on price of DVD issue.

**Gift cards – The final cut off date for ordering gift cards to reach you before Christmas day is 15th December. However, we recommend that you order your gift subscription before 1st December 2003 to guarantee your gift card will arrive in time. Orders can still be processed after 15th December, but gift cards are unlikely to reach you before Christmas day.

ORDER HOTLINE ☎ 0870 444 8645

LXF LUG pages proudly sponsored by



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

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 SCUNTHORPE & DONCASTER

URL www.scundog.org
Contact Shaun Holt – shaun@scundog.org

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.manlug.mcc.ac.uk
Contact John Heaton, Owen Le Blanc

25 HERTFORDSHIRE

URL www.herts.lug.org.uk
Contact Nicolas Pike

26 WEST YORKSHIRE

URL www.wylug.lug.org.uk
Contact Jim Jackson

27 SHEFFIELD

URL www.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.sderby.lug.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 Andy Hutchings A-Wing deltaone@virgin.net

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

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

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

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**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**URL www.euslug.lug.org.ukEmail 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 <http://sussex.lug.org.uk/>**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

88 ORMSKIRKEmail 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**URL www.aberdeen.lug.org.uk**94 SHETLAND**URL www.shetland.lug.org.ukEmail c_s_s_butler@yahoo.com**95 GLASTONBURY**URL www.glastonbury.lug.org.uk

Contact Steve Leonard-Clarke

96 SOUTHEND-ON-SEAURL www.sos.lug.org.uk

Contact Derek Shaw

97 ORPINGTONURL www.orpington.lug.org.uk

Contact Barry Schofield

YOUNG LINUXURL www.young.lug.org.uk**SCHOOLS**URL www.schools.lug.org.uk

LXF LUG pages proudly sponsored by



LUGS OF THE MONTH

Noosa Linux Users Group

Rick Phillips writes:

We are a very small LUG based around Noosa, a tourist area 150km north of State capital Brisbane. We meet weekly in rotation at my home and one of our member's business premises. No membership fees are charged.

Our interests vary widely from practical mutual help to just a "talk-fest". Some members have formed joint working relationships, but most importantly we enjoy each other's company. The group has been in

existence now for some six years and many members are the originals.

We are indeed fortunate in that the founder, Gerald Catling, has donated his dual CPU server and two other computers to the club for safekeeping while he and his wife Shirley travel around Australia for the next two years or so. Having these permanently available to the club allows members to conduct various experiments and to participate in training. Machines will be available to those who find it impractical

to bring their own to meetings. The server will also be set up as a 'distro server' storing copies of the latest ISOs for members' use.

We're very friendly here on the Sunshine Coast and we always like to hear from other Linux users near and far – here's our details:
Noosa Linux Users Group
c/o PO Box 97, Cooroy 4563, Queensland, Australia
Web: www.noosalug.org.au
email: rick@noosalug.org.au



Worldwide Linux User Groups

Free Software users across the globe

Africa

EGYPT
URL www.linux-egypt.org
GAUTENG, SOUTH AFRICA
URL www.glug.org.za
Email glugmin@revolution.org.za
THE LORD'S ABODE, JO'BURG, SA
Email Andrew Gargan avrin17@iname.com

Australia

ADELAIDE
URL www.linuxsa.org.au
Email mtippet@anu.edu.au
ALICE SPRINGS
URL www.aslug.org.au
MELBOURNE, VICTORIA
URL www.luv.asn.au
Contact luv-committee@luv.asn.au
PERTH
URL <http://plug.linux.org.au/>
SYDNEY
URL www.slug.org.au

Europe

COSTA DEL SOL (English speaking)
URL www.fuengirola.lug.org.uk
DENMARK
Alsund www.alslug.dk
Esbjerg www.eslug.dk
Fyns www.flug.dk
Midt-og Vestjylland www.mvjlug.dk
Nordjylland www.njlug.dk
Skåne Sjælland www.sslug.dk
Trekantsområdet www.tlug.dk
Vest-fyn www.haarby-net.dk/vflug
Århus www.aalug.dk
EIRE
URL www.linux.ie
Email root@linux.ie
URL www.dilu.org
Contact glossary@dilu.org
MILUG (Longford)
URL <http://midlands.linux.ie>
Contact midlands@linux.ie

Middle East

ISRAEL
URL www.iglu.org.il/IGLU/
Contact webmaster@iglu.org.il
PALESTINE
URL www.lugps.org
Email isam@planet.edu

Asia

HONG KONG (multilingual)
URL www.linux.org.hk
SINGAPORE – SLUG
URL www.lugs.org.sg
SRI LANKA
URL www.lklug.pdn.ac.lk
MYANMAR (formerly BURMA)
URL www.myanmarlug.org
Email afityde@balug.org
PAKISTAN
URL www.linuxpakistan.net
Email tux@cslug.org
HYDERABAD, SINDH, INDUS VALLEY
URL www.geocities.com/slug_pk/
KASHMIR
Coming soon!

China

BEIJING (GB encoding, but mostly written in Chinese)
URL <http://mud.263.net.cn/~linux>
CHINESE LINUX USER GROUP
URL www.linux.org.cn
NANJING
URL <http://jllib.jlonline.com/njlug>

India

LINUX INDIA
URL <http://linux-india.org>
ALIGARH LUG
URL <http://linux.amupost.com>
BOMBAY
URL www.ilug-bom.org.in
CHANDIGARH
URL www.geocities.com/vipinb
CHENNAI AND MADRAS
URL www.chennaiug.org/
CYBERABAD (CLUG)
URL <http://seeknew.freesevers.com/clug/>
DELHI
URL www.linux-delhi.org
KOLKATA
URL www.ilug-cal.org
MADURI
URL <http://linuxmadurai.tripod.com>
NORTHERN INDIA LINUX
URL <http://groups.yahoo.com/group/lug-northindia>

Spreading the word

"How can Linux make *my* business more profit?" Jono Bacon examines more pitfalls of advocating Linux use in enterprise.

One of the most difficult aspects of developing an advocacy campaign when targeting businesses is the sheer variation of business types, sizes and requirements. As you can imagine, the needs of a company like IBM are going to be a lot different to the requirements of 'Bob's Mobile Cheesemongers'. Although the size and type of business is different between Bob and IBM, some aspects will be shared. You know that both businesses will want a low-cost system, but obviously low-cost to a large multinational is probably not low cost to Bob and his part-time brother-in-law selling their fermented milk products, so the definition of low-cost needs to be specifically tuned for your target business. Another core requirement is support. Bob and IBM both need a person to point the finger at when something goes wrong. The issue of support is a stick which Microsoft is using to spread FUD about Linux, as you can see from our News Extra report on page 12

Support is a key area, but remember that support also depends on the

business. Bob may be happy with limited tech support over the phone or being trained up to make best use of any ad-hoc support that is available on the Internet, but an organisation that has even as few as half-a-dozen users may require dedicated support personnel on call, and the larger the business, the more on-site and more advanced helpdesk systems will be required.

Your job as an advocator is to find out how the size of the business relates to the way in which its IT is structured, and show an appropriate range of options available. Remember that cost is an essential issue, so if you can demonstrate in as non-technical terms as possible that they can get great software that works well, good support and better licensing for less cash, it is likely that they will pay attention to you. Any TCO (Total Cost of Ownership) figures should contain a margin for error – it's always impressive when a project completes under-budget!

Next month we will continue to analyse businesses and look at support in more detail. **LXF**

Linux User Group organisers

If you're not listed here, or we have your details wrong, please contact us at: **LUGS!, Linux Format, 30 Monmouth Street, Bath, BA1 2BW** or email your details to: linuxformat@futurenet.co.uk

HAPPY LXF-MAS!

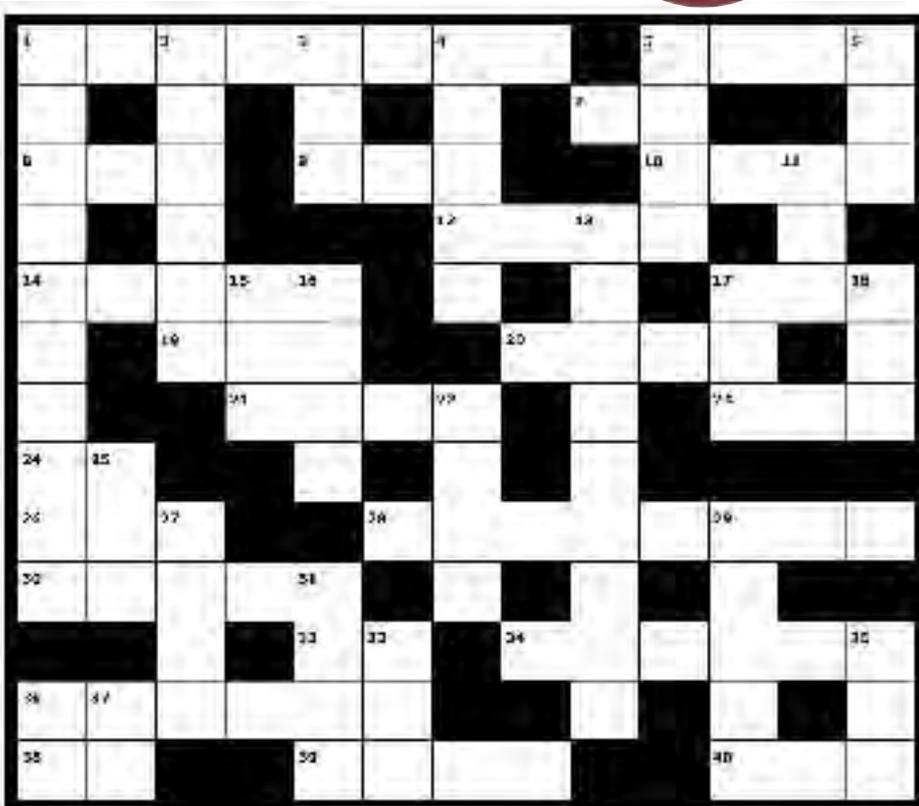


It's Boxing Day, and your family hasn't allowed you to turn your PC on for nearly 48 hours... Console yourself with some merriment from **Linux Format**.

CROSSWORD CLUES

Across

- 1 To his friends, he's Linus, to you he's Mr ____
- 5 Hear no evil, send it to /dev/ ____
- 7 Change us to a different user
- 8 Daemon that makes things happen at the right time
- 9 A hundred short of a changed deck gets you a nice desktop?
- 10 A shell inspired by maize perhaps?
- 12 Novell's latest purchase
- 14 Not just a text editor, more a way of life
- 17 A command that displays your current directory location
- 19 Mix the pot to find process performance stats
- 20 A language somewhere east of Krakatoa
- 21 Four-letter acronym used to describe software commonly deployed for webserving
- 23 A food additive or abbreviation for message
- 24 Command for underlining text
- 26 The most famous penguin in the world
- 28 Examples include *Opera*, *Mozilla*, *Konqueror*...
- 30 Used to search files for a pattern of characters, using full regular expressions
- 32 Abbreviation for the broadcast method used in wireless networks
- 34 SIGBRK is the signal to send when something _____
- 36 Change footwear to start again
- 38 Cold Fusions new suffix, or the country code for Mexico
- 39 The biggest repository of modules for 22 down.
- 40 command to create explicitly static symbolic links



Down

- 1 Show the path network packets take to reach a host
- 2 Thread together a Linux distro
- 3 Aho, Kernighan and Weinberger's initial processing language
- 4 Command used to display kernel boot messages
- 5 Post or PHP, an explosive web content system
- 6 Network smaller than a WAN
- 11 Make war unprocessed streams or data
- 13 The longest running Linux distro
- 15 When you need to filter reverse line feeds from input files...
- 16 _____, _____, eggs and _____
- 17 Authentication system one away from 16 down
- 18 You don't need a spade to lookup DNS info, just ____
- 22 A language that apparently came before swine
- 25 A handle, projection or organised group of Linux users
- 27 Utility used to get or set the contents of the RESOURCE_MANAGER property on the root window of screen 0
- 29 LPI's LPIC1-101 and LPIC1-102 are examples of these
- 31 Agricultural produce or a virtual directory containing live system info
- 33 Er, a protocol for file transfer
- 35 Evilly, a text editor or initial TCP packet
- 36 Delete a reversed title
- 37 Type of parrot, or command to start *vim* in a specific mode

ENTER... THE LXF MATRIX!

In case you've been living on Mars and not seen one of these before, the object of a wordsearch is to find the words hidden within the matrix grid. Words can be found horizontally, vertically, diagonally, backwards or forwards.

The wordsearch style of puzzle (call it whatever you like – WordSearch, WordSeek, WordFind, WonderWord, etc) was purportedly originally designed and Published by Norman E. Gibat in the *Selenby Digest* on March 1, 1968 in Norman, Oklahoma.

Selenby Digest was a small for sale/wanted

advert paper distributed free at Safeway and other stores in Norman.

The puzzle was very popular locally and several more followed this original. Some

```
L O O K W A N T L I N T
S T R I P A M R C E O U
D V D A P G G O F P X C
K R C M A K E F T B I S
D H I D E N T F S A N E
E G R E P H P T F S U Y
M I V N E M A C S H T M
O D W P R E D H A T A A
N A N A L E S S A P D F
G N U X N A L U S L F I
T A L A Y A Y S K I T N
K P L A S T S E D B U D
```

teachers in local schools asked for reprints to use in their classes, and the puzzle style rapidly grew across first the US, then the world, in newspaper syndications, puzzle books, and as teaching aides.

There's **over fifty** Linux-related words, abbreviations, commands and acronyms hidden in the above grid, some are obvious, some we didn't even know were there until we'd made the grid – exactly how many are there, and can you find them all?



BECOME A BLACK BELT... IN ORIGAMI!

Impress your friends and family over the seasonal holidays with your new-found mastery of the ancient art of paper-folding. Carefully follow the *LXF* step-through guide:

Step 1. From your basic 'fish' model, valley fold the bottom point of the reverse triangle down to the crease mark from 'Fish step 5'. This should look just like flap A on the other side. Unfold. Rotate the model anticlockwise 90 degrees. Valley-fold the top point up to where the crease lines meet. Next, reverse-fold the model in half. Make an outside reverse fold along the existing crease lines. Bring the left and right corners of the square together so that they tuck in at the bottom corner.

Step 2. Screw up your mis-shapen and heavily creased paper violently while muttering obscenities. Throw into nearest bin.

Step 3. When you miss the bin, curse loudly and kick both bin and paper across the room.

Alternatively, go to the following link for Frantisek Grebenicek of the Czech Republic's easily constructed origami penguin to sit proudly atop your monitor:

www.siskiyous.edu/computerlab/origami/penguin/

LIMERIX!

Some of the *LXF* team have wasted lots of their time coming up with some corny Linux limericks. If you think you can do any better, please send them in to lxf.letters@futurenet.co.uk, and we'll print a selection of the funniest non-libellous ones around Easter-time!

Three friends called *grep*, *sed* and *awk*
Set off to the park for a walk.
When *cat* said "boo hoo!"
They said "You can come too,
As long as you're sure not to fork!"

An aspiring young actress from Egypt
Asked her boyfriend to write her a script.
But he wrote it in shell
So she said "go to hell"
And cemented him up in a crypt.

There was a young lady from Leek
Who was horribly bored by a geek.
He told the poor girl
About Samba and Perl
And the programs he'd written last week

A Linux sysadmin called Pete
Typed all his commands with his feet.
After **rm -r**
His toe hit a star.
Now he's selling bananas in Crete.
Dr Chris Brown, LXF freelance writer

There was an old man from Tralee
Who wanted to revive his PC
So he installed the latest SUSE
Then boasted down the boozier
That Linux is the new Windows XP
Julian Jefferson, LXF Art Editor

There was a sysadmin named Jack
Whose server suffered an attack
He sshed in as root
And tried to reboot
But the system never came back

There was a young man named Nash
Whose attitude was rather rash
With no sense of alarm
He used each switch on *hdparm*
And paid with a nasty head crash
Nick Veitch, LXF Editor

There was a small company called SCO
That wanted our Linux to go
They sued all through the night
Without proving they're right
And soon they'll have no more dough!
Paul Hudson, LXF Reviews Editor

One night a penguin called Tux,
Said "I know why Microsoft sucks –
OSes and apps should be free
So why can't Bill see
That he stifles innovation with bucks?"

Mr Stallman is feeling quite blue
'Cause *LXF* leaves off his GNU
RMS claims "This is GPL robbery
That inflames Linux snobbery,"
But space-wise, what can we do?

When your servers always have to be on,
But Windows just keeps going wrong,
For maximum uptime,
Linux is sublime –
Soon your crashes are gone.

Linux packs functionality anatomical
Into hard disk space quite economical.
But so many distros I've tried
Leave my partitions all fried
And dependency hell's far from comical.
Matt Nailon, LXF Production Editor



The answers to *The Matrix* and crossword solution will be printed in the next issue of *Linux Format, LXF49, ON SALE 23 DECEMBER 2003.*



Essential disc info

Read this important information before you use your *Linux Format* coverdisc – CD or DVD. We've collated some helpful info to help you get the most from these jewels of data!

FINDING THE ESSENTIALS

Missing something?

As many of the programs on our discs are the very latest releases, they are often built on the very latest libraries and may depend on other packages your current Linux setup does not contain. We try to provide you with as many of these important supporting files and libraries as possible, though obviously we don't have space to include absolutely everything.

In many cases, the latest libraries and

other packages you might need will be included in the "essentials" folder on the disc, so if you are missing dependencies, this is the first place to look.

Package formats

Wherever possible, we try to include as many different types of package for an installation as possible, whether that be distribution specific RPMs, debs or whatever. Please bear in mind that we can only do this where space permits and when the packages are available.

We will, apart from exceptional or legally restricted situations, include the source files for any package, so that you can build it yourself.

Documentation

These pages provide helpful information on how to install and use some of the packages on the CD. Please note that many of the applications come with their own documentation, and there are additional notes and files in the relevant directories.

CREATING INSTALL CDS WITH CDRECORD

The quickest way to burn an ISO image to CD is with *cdrecord*. You need to be root to do this. First find the address of your CD-writer with

```
cdrecord --scanbus
```

This will show the devices connected to your system. The SCSI address of each device is the three numbers in the leftmost column, say 0,3,0. Now you can burn a CD with

```
cdrecord dev=0,3,0 -v /path/to/image.iso
```

You can simplify the command by saving some default settings in `/etc/default/cdrecord`. Add a line for each CD writer on your system (usually one) like this

```
Plextor= 0,3,0 12 16M
```

The first item is a label, after the SCSI address you put the speed and the buffer size to use. You can now replace the SCSI address in the command line with the label, but it gets even easier if you add

```
CDR_DEVICE=Plextor
```

Now you can burn an ISO image to disc with

```
cdrecord -v/path/to/image.iso
```

If you really don't want to use the command line, *gcombust* will do the job for you. Start it as root, select the "Burn" tab and the "ISO 9660 Image" gadget near the top of the window. Put the path to the image file in the gadget and press "Combust!". Now put on the kettle while the CD is created for you.

Other OS?

You don't have to use Linux to burn the ISO to a disc. All 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 has 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. 

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 disc being defective please email our support team (support@futurenet.co.uk) for further assistance. If you would prefer to talk to a member of our reader support team please call **01225 822 743**.

Coverdisc



Forget sherry and mince pies, Santa wants ALE! Neil Bothwick comes down your chimney with his bulging sackful of presents that is this month's *Linux Format* DVD.

DESKTOP KDE3.2

For a much more detailed explanation of KDE 3.2, please read our in-depth examination feature starting on page 50 of this very issue. Before we go any further, we should make it clear that KDE 3.2 as included on the *LXF* coverdiscs is an alpha release and not really suitable for use in a production environment. However, KDE 3.2 generated a great deal of interest at the recent UK Linux Expo, so in response to multitudinous requests from eager readers, here is a chance for more of you to see it in action.

This is a complete desktop environment and suite of programs, so installation is necessarily more involved

than for a single application. There are two ways to install KDE 3.2 from the provided source code tarballs. The standard method involves unpacking each archive and performing the standard call:

```
./configure && make && make install
```

If you choose to do it this way there are a couple of points to bear in mind. The first is the order of installation. You should install *kdelibs* first, followed by *kdebase*. These two packages are compulsory and should be installed in that order. After that you can install what you want. The other point is where you install it. To avoid overwriting your existing KDE setup, you should use the `--prefix` option to `./configure` to specify an alternate location. You can determine the current location with

```
kde-config --prefix
```

This is usually `/usr`, so you could install 3.2 to `/opt`. Make sure you use the same `--prefix` value for each package.

The less arduous installation method is to use *Konstruct*, also on the DVD. You need the unstable version for KDE 3.2 (the stable is for 3.1.4). *Konstruct* takes care of installing KDE 3.2 alpha in a separate directory.



KDE 3.2 is still in alpha testing, so be careful how you deploy it – and remember to ALWAYS backup your system before trying any new apps.

The default is `~/kde3.2-alpha2/` so you don't even need to run it as root. You should read the README file that accompanies *Konstruct* for details of how to use it. Make sure you set the `GARCHIVEDIR` option in the configuration file to point to the KDE3.2 directory on the DVD, otherwise it will try to download the files. KDE 3.2 needs a later version of Qt than you may have. If you find that

`./configure` complains, you can always get the latest version from <ftp://ftp.trolltech.com/qt/source/>, or alternatively let *Konstruct* download it automatically.

INTERNET MOZILLA

Development of *Mozilla* continues apace, despite the departure of AOL from the project. This month sees the release of new versions of three programs from the Mozilla.org project. *Mozilla* itself has reached version 1.5 and provides the best option if you want an all-in-one Internet package. This program is far more than a web browser – handling email, usenet news



Thunderbird is Mozilla's mail and news component, with some extras.

On the DVD

Wherever you see this logo it means there's related stuff on the DVD

IMPORTANT NOTICE

Before you even put the DVD in your drive, please make sure you read, understand and agree to the following: The *Linux Format* DVD is thoroughly tested for all known viruses, and is independently certified virus-free before duplication. We recommend that you always run a reliable and up-to-date virus-checker on ANY new software. While every care is taken in the selection, testing and installation of DVD software, Future Publishing can accept no responsibility for disruption and/or loss to your data or your computer system which may occur while using this disc, the programs or the data on it. You are strongly advised to have up-to-date, verified backups of all important files. Please read individual licences for usage terms.

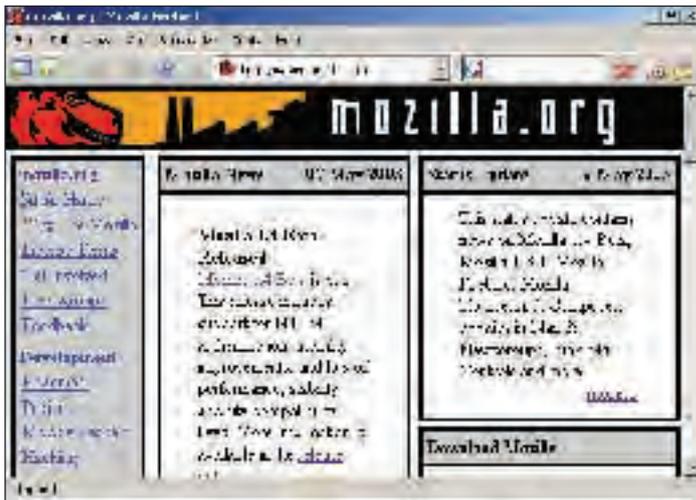
CAN'T BOOT, WON'T BOOT?

Incompatibilities with BIOS and isolinux

The first two CDs and the DVD are bootable, as is the Slackware CD created from the ISO image on the DVD. If you are unable to boot from them, the first thing to check is that your computer's BIOS is set to boot from CD before the hard disk. Your computer's or motherboard's manual should give details on this, but it usually involves pressing a key (often **Del** or **F2**) during startup to access a menu.

If your computer is correctly set up, but still refuses to boot from these discs, it is most likely that there is an incompatibility between your BIOS and

the isolinux system used by both the SUSE and Slackware bootloaders. The simple solution to this is to use a boot floppy. While each distribution CD usually provides boot floppy images, some of them require two or even more disks and it is a different procedure for each distro. You may find that all you need is *Smart Boot Manager*, in the Essentials directory of the discs each month. Copy the image from `sbootmgr.dsk` to a floppy disk, using `dd` from Linux or `rawrite` from Windows. Booting from this disk gives a menu that should enable you to select and boot from the CD or DVD.



Want to use *Mozilla* as a browser without the bloat of the extras you never use? *Firebird* may be just what you need.

and IRC as well as providing an HTML authoring environment. If you prefer the 'one program, one task' approach, the browser and mail components of *Mozilla* are also available in separate projects – *Firebird* and *Thunderbird*.

Firebird is basically *Mozilla* without the extras, concentrating on being a web browser and nothing else.

Thunderbird is a mail/news reader and offer a wide range of features for heavy mail users, while still being easy to set up and use for those with lighter needs. *Thunderbird* is not just *Mozilla* without the browser, it now has more features than the program it came from. *Firebird* is at version 0.7 and *Thunderbird* at 0.3, so these are still at an early stage compared with *Mozilla* itself, although they reap the benefits of using the well tested code from the parent project.

This month's discs are something of a browser bonanza. We also have

Epiphany, another browser using Mozilla's Gecko HTML engine; and *Opera 7.21*. This is an important release of *Opera* from the Linux perspective as it is the first time *Opera* has released the same version for Linux and Windows at the same time.

SERVER MAMBO SITE SERVER

There was a time when creating a website meant hand crafting your own static HTML with something like *Emacs*; while the more user-friendly alternative was to use *Vi*. While a good understanding of how HTML works is never a bad thing (see LXF's new tutorial series beginning on page 66), it never hurts to let the computer take over the laborious aspects of the task. Some may say that a content management system (CMS) goes too far, but it does allow one person to administer a site to which

PARTITIONING YOUR HARD DISK

More than one way to root and swap



A hard disk is divided into partitions, each of which is accessed as a separate data storage area. A standard Windows setup has a single partition occupying the whole drive, so you may not have even noticed that it was partitioned before, but adding a second operating system will require that each has its own portion of the disk. Unlike Windows, Linux installations generally use two or more partitions, the exact number and layout depending on what the machine is being used for.

Ask ten Linux users how you should partition your hard disk and you will probably get ten different answers. The simplest setup has a swap partition and a root partition. The swap partition should be roughly twice the size of the

computer's RAM, but probably not more than 512MB. It is used for temporary storage of data to free up memory. The root partition (/) contains everything else.

You can also have separate partitions for other sections of the filesystem hierarchy, /usr, /var/, /home and /usr/local are the most common candidates for this. However, this normally only complicates the situation. It is generally worth having /home on its own partition. This contains your own data and configuration files, making it a separate partition means you won't lose it if you reinstall. It also means you can share one /home partition between more than one distribution, if you like to experiment with the different distributions on the LXF coverdiscs.

many contribute. The administrator should still have a good working knowledge of the underlying technologies, but those contributing the content should be free to concentrate on the quality of that content.

Mambo is one such CMS. Whether you are creating a small site or one of

thousands of pages with a large number of contributors, it could save time, and that time can then be spent on making the site more useful. There are two tarballs on the disc, the current stable version and a beta of the next release.

There are also a number of components, modules and templates that can be used to extend the functionality of a site or to change its appearance. If you like the look of *Mambo*, you should visit the project's own website. This includes a number of tutorials on using and creating *Mambo* components to enhance both the appearance and performance of your website. [LXF](#)

INSTALLING SLACKWARE WITHOUT A DVD DRIVE

Building ISO images

As is now normal with our bootable DVDs, we have provided a way to create CD ISO images so that you may burn your own CDs for installation on a computer without a DVD drive. This can be done from Linux or Windows, although there is a problem creating the second disc from Windows this time.

In order to build the ISO images in Linux, type one of the following commands into a terminal:
`sh /mnt/cdrom/Distros/Slackware/mkiso`
 or
`sh /mnt/cdrom/Distros/Slackware/mkiso /path/to/some/dir`

The latter creates the ISO images in the specified directory, the former in the

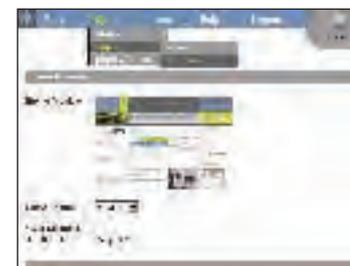
current directory. You can also run the script by clicking its icon, although I have only tested this in KDE, when the ISO images will be put in your home directory. Naturally, if your DVD drive is mounted somewhere other than /mnt/cdrom, you should make the appropriate change.

Windows users can create the first CD by double-clicking the *winmkiso* icon or running the script from a MSDOS prompt. In the latter case, you should change to the *Slackware* directory of the CD before running the script, and you can give an optional destination as an argument. With no argument, or when run from the icon, it puts the ISO images

in C:. For example, if your DVD drive is E: and you want to save the ISO images to D:ISO

```
cd Distros/Slackware
winmkiso D:ISO
```

The script may fail to create the second CD – it has done on some occasions when tested on a Win98SE system. This is because the CD filesystem used by Windows was unable to read some of the filenames on the DVD. However, you only need the first CD to install a working system, you can copy any extra packages you need from the DVD to a CD-R and install them later.



Create and manage a website with the *Mambo* CMS.

Welcome

Twenty pages of real-world Linux for IT professionals

There has been a lot of scaremongering about SCO's legal action, and the potential outcomes. Some of this has obviously been whipped up by SCO itself, and other shadowy individuals and organisations with an axe to grind on the subject of Free Software. The scaremongering has certainly led to questions being asked by potential users, but thankfully hasn't. It's not about 'Free software hippies' taking on big business. If SCO was to get its way, it won't be the Linux developers or users that lose out, it will be business and academia. SCO isn't going to knock on the door in the middle of the afternoon and rouse the bedroom coder from his slumber; they will continue to target the people with the most money. And that isn't just a threat to those companies who use Linux, but to everyone who uses a computer.

Thankfully though, it seems unlikely that there is much of a case to answer, judging by the amusing evidence already produced by SCO. The latest rumours are that SCO's legal team intend to try and prove that the GPL is 'unconstitutional'. It strikes me as rather odd that the forefathers of the 'greatest democracy the world has ever known' would have wanted to make it unlawful to give things away. But then I'm not a very expensive lawyer, so perhaps I'm missing some subtle nuance. If you want to find out more about the potential ramifications and – win or lose – what the Linux community can learn from the whole affair, I direct you to our feature starting on page 6.

Meanwhile the world still turns, and there are plenty of other things to keep abreast of. For a start, the popular JBoss application server now has extended support options in Europe (page 14), Novell has reinvented itself and enterprise Linux (page 16), and of course, we still have our regulars on security (page 4) and another Linux hosting profile (page 18) to enjoy. Let me be the first to wish you all a Happy New Year!

Nick Veitch Editor
nick.veitch@futurenet.co.uk



“It strikes me as odd that the forefathers of ‘the greatest democracy in the world’ would have wanted to make it unlawful to give things away for free...”

CONTACTS

EDITORIAL

Editor

Nick Veitch nick.veitch@futurenet.co.uk

Reviews editor

Paul Hudson paul.hudson@futurenet.co.uk

Art editor

Julian Jefferson julian.jefferson@futurenet.co.uk

Production editor

Matt Nailon matthew.nailon@futurenet.co.uk

Publisher

Kelley Corten

Group Publisher

Dave Taylor

Photography

Photodisc, Corbis

ADVERTISING

Senior Sales Executive

Andrew Hudson

andrew.hudson@futurenet.co.uk

(01225) 442244 Ext 2441

Copyright No part of this publication may be reproduced without written permission from our publisher. All letters & 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. *Linux Format* recognises all copyrights in this issue. Where possible, we have acknowledged the copyright holder. Contact us if we haven't credited your copyright and we will always correct any oversight. We cannot be held responsible for mistakes or misprints.

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.

LinuxPro is a part of



magazine

Published by



Media with passion

30 Monmouth Street, Bath BA1 2BW
Tel +44(0) 1225 442244

Contents

SECURITY

ISS's CTO highlights the biggest challenges facing sysadmins **p04**

SCO VS LINUX

IP barrister David Harris assesses the legal implications **p06**

JAVA DEVELOPMENT

JBoss marches into Europe **p14**

ENTERPRISE LINUX

Novell commits to Open Source **p16**

HOSTING

Host Europe's across-the-board solutions for SMEs and beyond **p18**



Serious about security

While IT workers have become more aware of their role in security since big viruses like Melissa struck some years ago, it's still clear to see that as protection technology has advanced, so has the software used by hackers.

Technologically speaking, we're less secure than ever before, so why do so few companies take a serious enough stance?

Linux Pro spoke to Chris Klaus, the Chief Technology Officer and founder of Internet Security Systems (ISS) to find out where things are now, where they will be in the future, and how ISS is making sure that companies are as protected as they can be...

LINUX PRO: Security is a big field in computing, and it's always growing. Where does ISS fit in?

CHRIS KLAUS: ISS is a world-leader in products and services that protect critical information assets from an ever-changing spectrum of threats and misuse. Our products dynamically detect, prevent, and respond to sophisticated threats to networks, servers, and desktops. They include *Internet Scanner* for vulnerability management and the *RealSecure* and *Proventia* range for intrusion prevention. Services include 24/7 managed system monitoring, emergency response, and access to the X-Force™ – the renowned ISS research and development team.

ISS is the trusted security provider for more than 11,000 corporate customers, including all of the Fortune 50, the top 10 largest US security firms, 10 of the world largest telecommunications companies and major agencies and departments within US local, state and federal governments. ISS is based in Atlanta, Georgia, and has offices throughout the Americas, Asia, Australia, Europe and the Middle East.

LXP: What are the biggest challenges sysadmins face?

CK: In large organisations the change management process is very cumbersome. It is sometimes difficult to figure out who owns the vulnerable assets; which cost centre should bear the expense; or even who is responsible for the security patching. Additionally, many production servers only have a small window of possible downtime (like two hours a month) and it becomes very difficult given that updates are coming out weekly.

System administrators are faced with three major challenges in security:

- Security sysadmins are having to learn to view security as a business risk and have to be able to discuss it in business terms, mainly in order to justify getting sufficient budget for their projects. Security advisors like ISS are helping the sysadmins put together financial risk models for businesses in order to calculate the return on investment. With enough

When it comes to locking hackers out of your network, who are you gonna call? PAUL HUDSON has been speaking to the CTO of one of the world's leading IT security companies...

data now available from many third party sources like the FBI/CSI and insurance companies on the impact of not being protected, it is feasible to calculate the probability of risk to a particular business and what it will cost them to get protected.

- The ever-growing issue of staying patched against the latest vulnerabilities. The number of vulnerabilities is increasing at a faster rate than security budgets. Patching multiple systems in a large environment is problematic and hard to scale. Many companies lack vulnerability assessment tools, therefore it is hard to determine what has, and has not been patched. Manually patching provides no visibility, nor reports of progress and trend lines. There are of course dependency issues too – for example, patching a web server may break applications running on it. It is extremely difficult to patch any legacy applications and operating systems since they are often no longer supported, plus most security patches force the system to reboot, which can be very disruptive to a production server.

- There are more and more point security solutions available and each niche product represents additional cost in management. Today we have firewalls, VPNs, anti-virus, Intrusion Detection Systems (IDS)/IPS (Intrusion Protection Systems), anti-DoS (Denial of Service), anti-spam, content filtering, etc. Every time a new point solution appears, it has its own separate management console, with its own policy, configuration, deployment, events, correlation, reporting and updating. Interestingly, all of the point solutions do the same thing with network traffic. The security industry will see large amounts of convergence happening over the next 6 to 12 months, driven mostly by cost-effectiveness.

Using protection agents on Linux hosts, like those from ISS with intrusion prevention and Virtual Patch(TM) technology, will act as a front line of defence and can reduce the risk of compromise in a similar way to security patching, but at far less cost. Updating Virtual Patch via the protection agents is much more scalable and has less impact than updating using security patches. More companies are beginning to use protection technologies as their first line of defence and using security patching as a defence-in-depth strategy.

LXP: What are the security pitfalls people most commonly fall into?

CK: The biggest security pitfall is the lack of awareness of an organisation's security baseline.

Most organisations do not have any idea what, and where the security holes are in their infrastructure, or on which mission-critical assets. There are many security assessment

tools for Linux that give a good understanding of the security profile of the network. Understanding the current security snapshot of the hosts and network provides a good direction for addressing the security issues.

There is a pitfall that many people fall into attempting to get their systems protected. Their approach is to try to patch and lock down the systems after they have been deployed. This is both too reactive and costly. A more proactive approach is to build a hard drive image where the OS has been patched, locked down and hardened. They can use this image as the template for all new desktops or servers. This allows the organisation to deploy many new desktops and servers in a secure mode while opening them up just enough to allow them to do their business.

LXP: To what extent does security factor into the business plan of today's companies? Is it changing?

CK: Not enough – it's usually an afterthought. There are many regulations and standards, like the UK's Turnbull Review (in late 2002 it became obligatory for listed companies to give an evaluation of the likelihood and impact of a wide range of risks becoming a reality in their annual reports) and the USA's HIPAA, Sarbanes-Oxley and Graham-Leech Bliley Acts. These are making security a 'must have', versus a 'nice to have'. Security will need to be built into more business plans.

LXP: What do you think motivates crackers to attack, control, and damage systems?

CK: The cyber-crime world is a reflection of the physical world. The motives vary depending on the cracker or criminal. They can be doing it for fun, for ego, for making a political or religious statement, or simply to make money. The major difference between physical and Internet crime is that cyber-crime requires some computer skills, but a keyboard's a lot less dangerous than a gun. Also, the chance of getting caught in the cyber-world is less, while the rewards are greater. Cyber-crime is becoming a growth industry as businesses and society become dependant on the Internet.

LXP: What do you feel is the most important single step to keeping a network secure?

CK: Simple – understanding the risk and applying the appropriate security solutions.

Most companies do not have a good assessment of their security profile. They really need a third party to provide a security assessment and penetration test against their mission-critical assets. Sometimes the CEO and board may

OPEN vs CLOSED

LXP: DOES OPEN SOURCE HAVE THE POTENTIAL TO BE MORE OR LESS SECURE THAN ITS CLOSED-SOURCE EQUIVALENT?

CK: Both Open and closed-source have many security holes. The vulnerabilities do not differentiate between Open Source or closed-source. It is really down to the developers behind each, how secure they are.

Many of the vulnerability issues like buffer overflows are well known in security circles, but not all developers and programmers know about security issues, nor do they know how to identify them. Therefore, these types of programming bugs get reintroduced in many applications and programs regardless of whether they use Open Source or closed-source. Most security consultants now have tools that identify these security vulnerabilities in either platform, so there is then little advantage of one over the other from a security point of view.

What will make the source code more secure (or less) will be whether developers and programmers understand the security implications of their designs and how they write their programs, along with a good security QA (Quality Assurance) process.

pay attention better to an independent body than their own security team. Without buy-in from upper management, it is hard to get budget and appropriate resources to address the security issues. All employees, from the top to the bottom, must understand the risk in terms of how it affects them and what they can do to reduce it. Starting with a security assessment to understand the organisation's risk is the first big step in moving the ball forward.

LXP: Where do you think IT security is going in the next few years?

CK: IT security is going to evolve into a security platform that provides comprehensive protection against all threats, including Denial of Service, hackers, worms, viruses, spam, etc. There will be convergence into security platforms with fewer stand-alone security management systems. All the different security solutions must start to work together and appear seamless to the security operator.

Many gateway network appliances will be delivered on Linux OS. Linux will be a key component in protecting the entire infrastructure, including all hosts and routers. More Linux protection agents will need to be deployed as well.

Security must move from a reactive to a proactive mode in the next few years. Too many organisations and security solutions react after the hacker, worm or virus is attacking their assets. When new viruses can spread around the Internet within fifteen minutes, like the SQL Slammer did, a security solution that is updated a few hours later is too late. New security solutions that proactively stop the viruses and worms need to be deployed.

Most organisations today do not have 24/7/365 security monitoring, especially during weekends and holidays, when hackers are most active. Their security staff are not resourced to provide this kind of functionality. Therefore, more on-going security management and monitoring is going to be outsourced over the next few years. As companies look at the business risk and ask how to do it cost-effectively, the outsource model makes the most economic sense. In the physical security world, most banks do not directly hire the physical security guards nor watch the alarms, but rather use the guards and alarm monitoring from a security company. Many organisations are making the same decision for the cyber-security world. There is just not enough security expertise – it is expensive, and security is not usually core to most businesses. While security is extremely important, it is not why most businesses are in business. Businesses should focus on what they are good at.

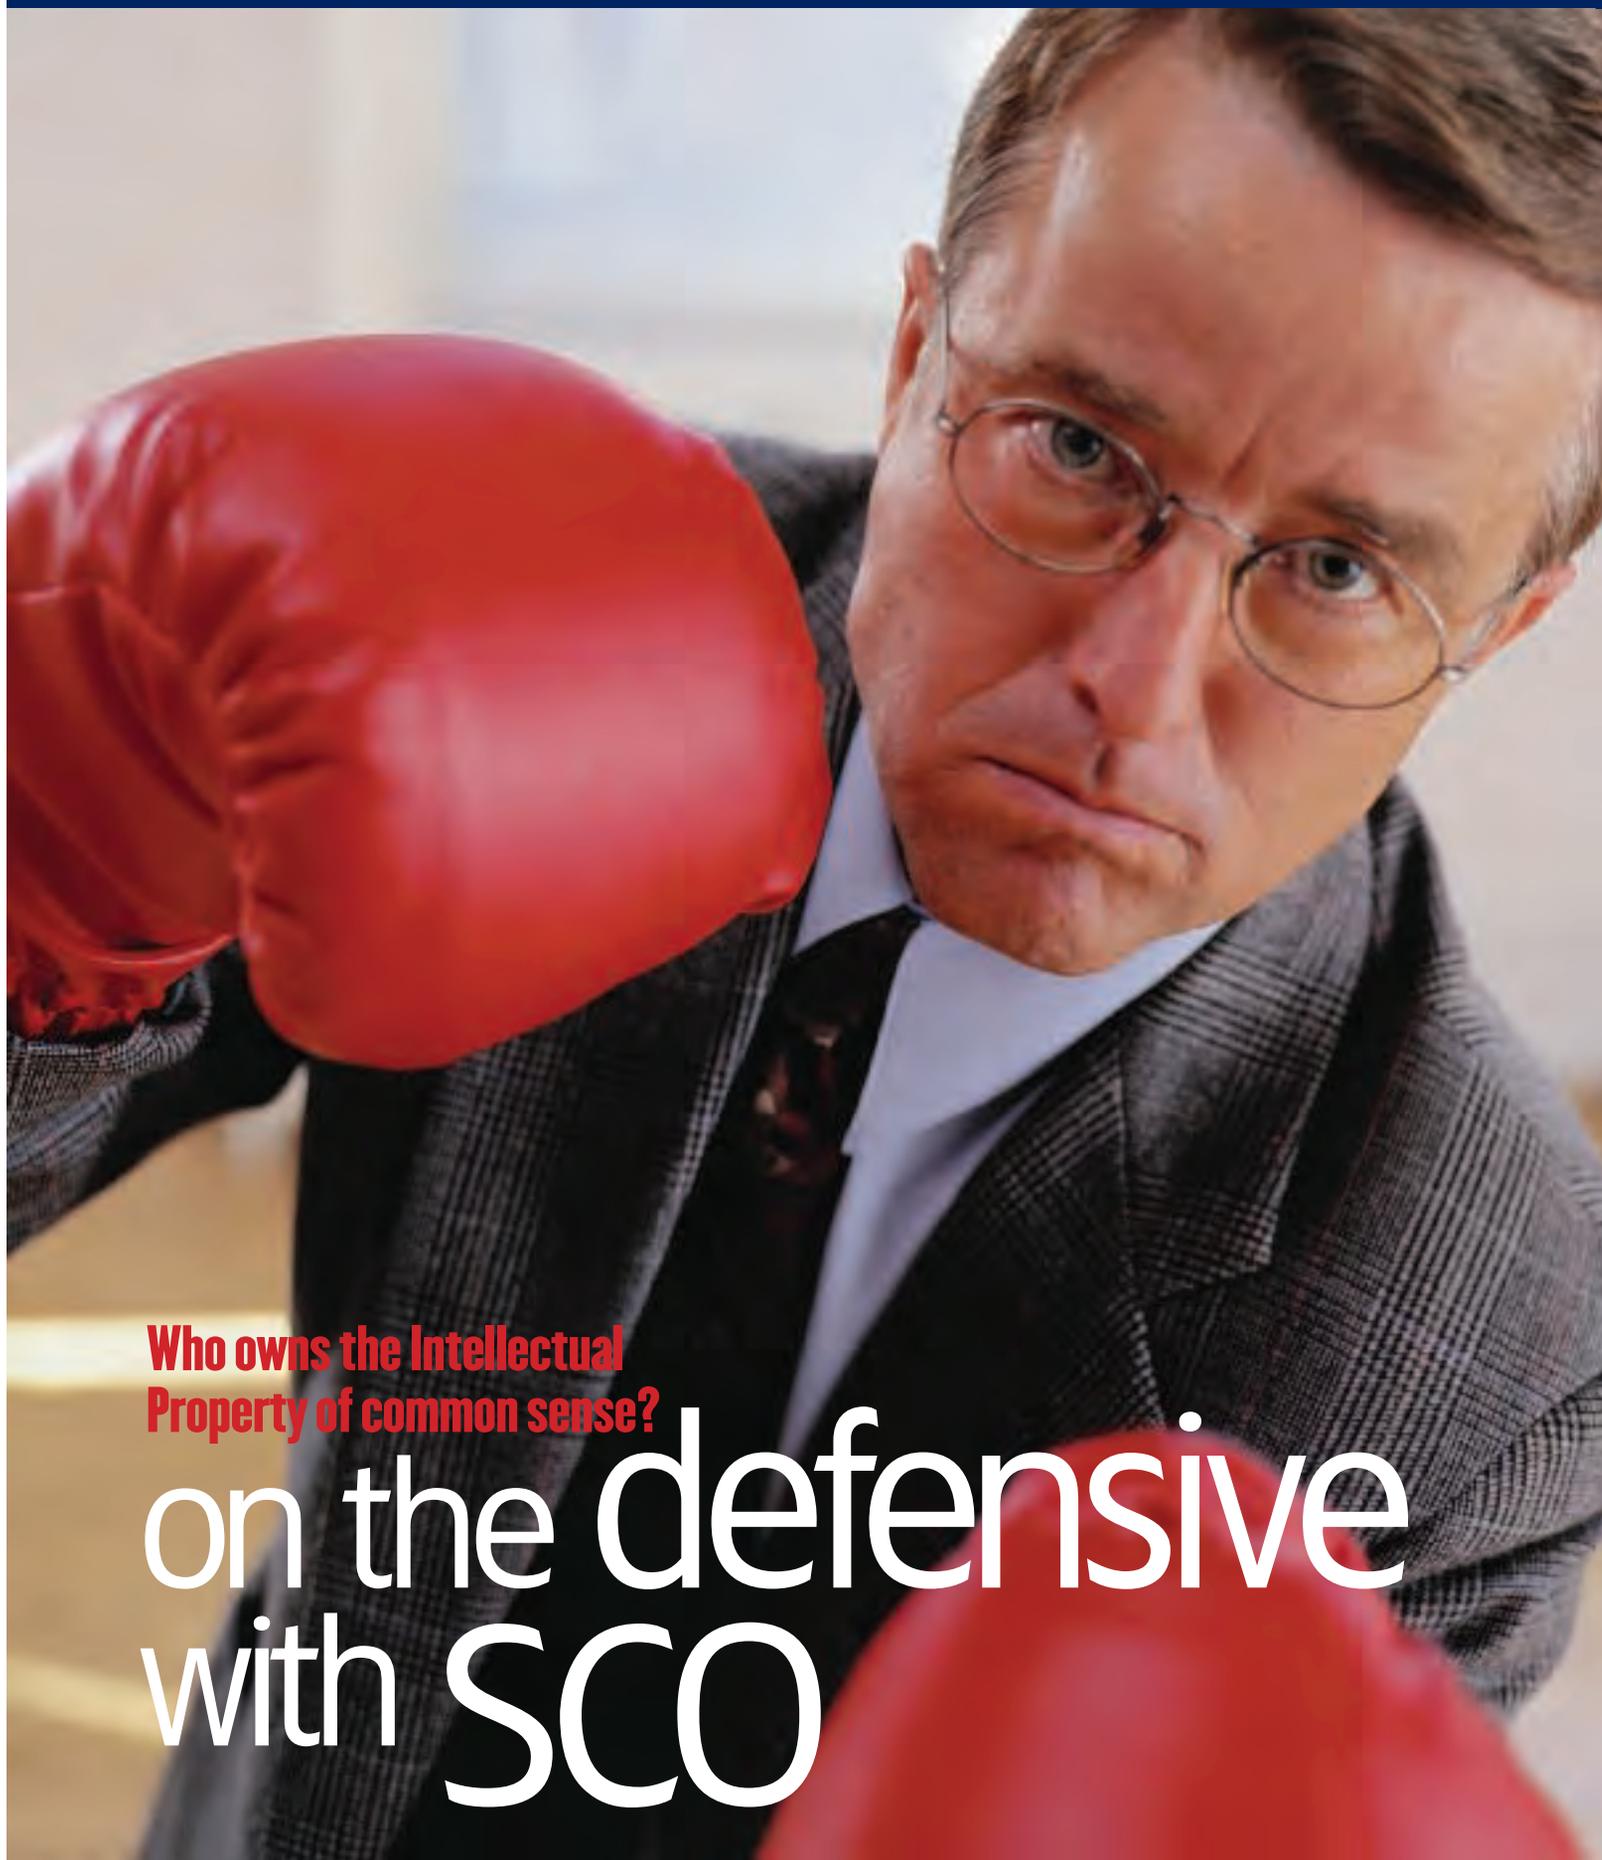
Does this mean that all of security will be outsourced? No! The business will still need a security staff to develop security policies, standards and build security into their infrastructure designs. But the mundane security operations of monitoring for alerts and alarms and managing the actual security devices like firewalls and IDS are much more effectively done by an outsourced security partner.

Organisations are starting to realise that cyber-security is an important business risk that needs to be dealt with, and are finally starting to spend money on security, which helps foster more security research and advancements, which in turn will help tip the scales in favour of the businesses against the criminals. ■

“Join the full-disclosure security mailing lists, and work with a trusted security partner to keep yourself ahead of threats.”

CHRIS KLAUS, CTO at ISS





**Who owns the Intellectual
Property of common sense?**

on the defensive with SCO

The SCO vs IBM case is unlikely to damage Linux, but as **DAVID HARRIS** explains, as well as being handy publicity for alternatives to restrictive copyright, it does highlight the need for careful coding practices and some of the risks that must be considered by FOSS developers.

Well, SCO has certainly put the cat among the pigeons. Back in issue 40 we discussed liability in general terms; peer-to-peer; criminal liability and so on. *SCO v IBM* had just kicked off and was looking like an important – albeit not central – issue in Free software. Since then, IBM has refused to either buy SCO or pay it off, causing the company to have a temper tantrum and throw its rattle out of its pram. SCO has now decided that Linux is a hippie pot-smoking communist conspiracy to steal its code and undermine the American way.

SCO has retained well-known lawyer David Boies, famous for losing against Microsoft, losing for Napster and losing for Al Gore; Linux fans can only hope his performance remains consistent. It is still way too early to be certain of the outcome of the case, regardless of the apparent inadequacies of SCO's position so far. Nonetheless, the whole debacle has raised more uncertainty among some commercial users of Linux than would appear from the outward expressions of faith and confidence among the user community.

Money talks

Who would be best served by such uncertainty? Well, SCO would if it caused people to buy their licence at US\$700 a pop, but thus far there is virtually no sign of them doing so. However, there is someone else lurking under the eaves: our old friend Microsoft. It may well be that despite paying \$10,000,000 in supposed license fees, they are entirely uninterested in the outcome of the case. Equally, since Microsoft uses relatively little Linux, the view is that they are seeking to wage war by proxy. If SCO succeeds in its lawsuit, its next attack is likely to be on the 2.2 & 2.3 kernels, and eventually Linux becomes a hostage under their control. SCO would then be bought out, marginalised or undermined by anti-competitive pricing by Microsoft, and the threat of Linux would finally be neutered. At worst, Microsoft buys time while the Free software community tries to replace Linux with an alternative such as a fully clean Linux or HURD.

What is more important – and a greater threat – is that the core of Free software risks being tainted. For several years, Microsoft has waged a campaign about the risks of Free software to intellectual property. It says that Free Software will contaminate proprietary products and make IP ownership uncertain; or it says that the use of Free software will cause proprietary rights to leak into the public domain via the GPL. With the SCO case, MS has a real example to point to and say “Ah-ha! See, we told you. If you dabble in the GPL this is what will happen. You risk both being sued and having

SCO'S COMPLAINT

SCO claims to own the intellectual property of UNIX

as a result of its purchase in 1995 from Novell of the original AT&T codebase and patents (see timeline *SCO vs Linux* on page 11).

This is however a deceptive aspect of SCO's case: it has claimed in several statements to own the UNIX operating system, and through numerous contractual arrangements to be entitled to control the rights of all vendors to use and distribute UNIX. At the same time it has acknowledged that

DISCLAIMER

THIS ARTICLE IS BASED ON UK law except where otherwise indicated. Substantial differences exist between UK law and that elsewhere. The consequence of all law varies greatly with individual circumstances and thus nothing in this article is intended as or should be construed as advice or acted on without seeking your own lawyers' advice. The author regrets that he cannot give personal legal advice, except in very specific cases, as is stated in outline at the bottom of page 12 of this issue.

your IP stolen”. As we all know, society will collapse if Microsoft can no longer make monopoly profits and its hegemony of the IT industry is threatened...

What has happened is that a change of management at SCO has led to a re-evaluation of profitability, resulting in a view that cannot see revenue potential from FOSS – possibly since the competition is too great and its own offerings are too poor. The intent is to revert to older business models, and so SCO needs to renege on its deal with the FOSS community. At the same time, SCO needs to eliminate the competition that previously devastated them by making highly questionable claims about Linux IP. Dubious claims of copyright infringement appear to be a part of Microsoft's attempt to effect that revisionism.

Clearly, the potential ramifications of SCO are serious and profound, and may prove damaging to FOSS in ways as yet unclear, but at this point I shall remove my tin-foil hat and get down to the legal analysis and commentary.

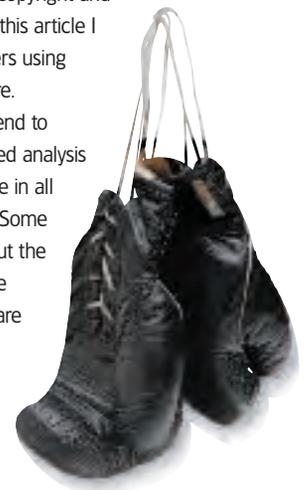
Several companies have asked me questions on the SCO case, like:

- Does SCO have good prospects of winning?
- Should they stall implementation while SCO goes on?
- What can be done if SCO do win?
- What are the implications of mixing of proprietary products and the GPL?

“SCO has retained well-known lawyer David Boies, famous for losing against MS, losing for Napster and losing for Al Gore.”

A number of common themes in copyright and licensing arise in these questions; in this article I hope to guide companies and hackers using contributing and writing Free software.

First, I should say that I do not intend to provide an exhaustive or even detailed analysis of SCO's claims in their entirety, since in all truth, they vary from week to week. Some cynics have pointedly remarked about the regularity of announcements and the consequent see-sawing of SCO's share prices, but we don't want to go into that. No, I intend to examine a few SCO complaints that illustrate more general problems in FOSS.



the Open Group is the owner of the UNIX trademark. The Open Group has however stated that *UNIX* is not a particular code implementation but rather a trademark and an associated specification. The Open Group certifies certain implementations of the standard as being conformant to it and hence being *UNIX*. There can thus be multiple Unices; this has already happened with IBM's OS/390 which has been certified as UNIX 95 and thus *UNIX*. Clearly, SCO does not own UNIX (the trademark) but one

COVER FEATURE **SCO**

particular historical codebase. As we discuss below, it may not even own all of that. Further, SCO claims to be able to licence further use of UNIX but this appears not to be the case: UNIX is a trademark that they do not own, but that they merely have a *licence to use*. SCO has no power to sub-licence the trademark from the Open Group, and hence on one argument, it has no power to grant rights to derivatives of UNIX. There is a countervailing argument that they do have an entitlement to grant a licence to the UNIX trademark if their original contract with the Open Group is to have any commercial meaning. However, even if true this would not give them the *exclusive* power to licence anyone else's UNIX implementation. Someone such as Linus Torvalds, the FSF or a wealthy benefactor could pay the Open Group to certify Linux as UNIX-complaint.

“This isn't a complex case – it's just that the facts have been hidden by SCO, and as a result, the case may be long and costly.”

Of course, none of this is relevant to the issue of code copying. If SCO's version of events is correct, and those 'Linux hippies' have stolen SCO's valuable IP and added it to the Linux kernel, the above arguments would be an irrelevant side-issue. It does however seem a little surprising that SCO, who have hired a supposedly top-notch team of litigators, haven't fully debugged their case statement and seem to be indulging in some generally sloppy work.



Copyright infringement

Direct copying

Is there any real basis for SCO's claims of copyright infringement? SCO seems to think so, but the evidence so far has not been entirely supportive of the thrust of its claims. During its resellers show, to support its vague claims of millions of lines of stolen code in the kernel, SCO revealed a sample code fragment partially obfuscated in a Greek font. This was a mistake, since hackers decrypted it rapidly and demonstrated it to be code whose heritage went back as far as 1968 to Donald Knuth. The code in question bears an SGI copyright notice, but SGI seems to have been somewhat cavalier with copyrights; it appears to have stripped the old copyright notice and replaced it with its own. This was clearly wrong and a copyright violation on SGI's part. This is in breach of both the advertising clause of BSD and the non-removal of notice clause. This file has been removed from the kernel, but that doesn't alter the malfeasance in SGI's act. Simultaneously, there is yet another version of this same code fragment obtained elsewhere under a BSD licence that would be perfectly legal to put in the kernel.

Taken at its worst, this infraction would result in SGI facing damages of a few tens of thousand of dollars, and probably much less, for altering attribution and what was a technical infringement of the licence; after all the code dissemination in that file had been agreed to by SCO (or its forbears). Would this merit a Linux licence fee of several hundred dollars a CPU to SCO? Not in any universe near me.

The other code fragment SCO laid claim to at the show is the *Berkley Packet Filter (BPF)*. This has been released on a BSD licence that Linux and SCO have always been able to use, and it seems that SCO copied it into System V. The Linux implementation of *BPF* however is an original work, based on a specification not owned by SCO. When SCO points to code similarity, it may well be not so much that Linux has copied SCO code, but that both have the same parents.

Clearly, this indicates how complex the factual matrices in litigation can arise, and this is not really a complex case – it's just that the facts have been hidden by SCO, and as a result, the case may be lengthy and costly. One can speculate however; as Bruce Perens has said, it is plausible that SCO have chosen the best examples they have and these are they. If so, then the danger to Linux seems minimal. None of this would justify the removal of 2.4++ kernels under an injunction nor any substantial damages. One hopes that any other examples are no more damaging. The apparently aberrant behaviour of at least one SGI employee does raise serious questions about code control, and is one of the more reasonable points made by SCO that should be absorbed, as I discuss later.

SCO's claims of infringement rest at two levels. One is the literal copying we have just discussed. We have seen several published examples so far; however, the code copied is either in the public domain, or the copying and derivation of the code by Linux developers has been agreed to at a conceptual level – it is merely that SGI took credit where they were not entitled to. It is not as though vast tracts of source code that had cost SCO a fortune to develop and had laboured hard to keep secret were stolen by industrial spies. We shall see in the fullness of time

whether SCO actually has better examples of copying. Nonetheless, copyright infringement is wrong, it should not happen and the FOSS community needs to be continually aware of the risk and make efforts to avoid it.

We can, I think, be less diffident about the other claim of infringement. The *Berkley Packet Filter* claim seems to be plain wrong on the facts so far disclosed. *BPF* was based on code written by Carnegie Mellon University put into 4.3 BSD and then copied by SCO into System V. They were entitled to do that, but not to claim, as they are now doing, that they own it and can eject others from its use.

Non literal copying

Not all copying is line-by-line, sometimes there is a deliberate attempt at obfuscation by infringers attempting to cover their tracks. At its crudest, this may involve no more than stripping off copyright notices changing variable names and shuffling lines of code around. This is crude, not uncommon and effective enough in most cases. Equally however, there can be a copying of an idea and significant chunks of its expression in reworked form. When this type of case comes to court, the assertion of copying is tested by first asking whether there is any causal connection; did the defendant copy it from the plaintiff, did the plaintiff copy it from the defendant or did they both copy it from some common source? If this sounds familiar in the context of the SCO complaint so it should. To determine such causation, questions are asked such as 'how similar is the code' 'were the *dramatis personae* linked in some way?' - eg was the alleged copier an ex-employee of the plaintiff? 'Was there an opportunity to copy the code? Did the plaintiff show it to the defendant for some reason?' In the SCO case, the broad accusation against the FOSS process is that various people in companies with access to System V source have, in solving coding problems, peeked at the System V source and obfuscated it into Linux.

Subconscious copying

Another form of indirect copying often seen is subconscious copying. It may well be that the defendant believes - in all good faith - that what they produced was their own, but in reality they have seen another work, forgotten it and then reproduced it. Most of the cases have come from the music industry for example George Harrison subconsciously copied parts of the tune for *My Sweet Lord*. This form of copying is less frequently seen with software, possibly because the more sophisticated subject matter means the scope for substantial taking (see below) is reduced.

Core protect-able ideas

The other main form of non-literal copyright infringement is indirect copying, and it seems that SCO originally seemed to be accusing Linux of this. If it were not possible to infringe copyright by copying a copy, the practical effectiveness of copyright law would be very damaged. One could turn Free software into a proprietary product in defiance of GPL merely by making a copy of someone else's infringing source code. Not a good thing.

A simplistic truism has developed in the FOSS community that one can only protect the expression of an



Darl McBride, SCO's CEO - stock-pumping son of Satan or ultra-shrewd businessman, depending on exactly who in the computing world you ask.

idea rather than the idea. To paraphrase one judge: "yes, but it rather depends what you mean by 'expression'". The point being that there is seldom a fixed dividing line between an idea and a fully developed exposition of an idea. True, no one can have a copyright in detective novels, but equally one cannot write a Sherlock Holmes novel with a hound and a moor. What about an English detective who employs forensic techniques and has a male sidekick? The point here is that in reality, there is a sliding scale between an idea and a its concrete expression, and whether there is infringement depends on where on that line you fall. This is referred to as *substantial* taking, and at varying stages appeared to be a SCO claim: that they owned System V and all Unix IP. SCO says that by copying core concepts, using the inspiration of Unix, or subconsciously copying Unix as result of seeing it at college (euphemistically called 'learning'), there is infringement. In assessing alleged copyright infringement, the principles we've mentioned would be used to assess whether there had been substantial taking. In the US, when software copying is alleged the complicated 'Abstraction Filtration comparison' test, from *Computer Associates v Altai* is used. This involves a three-step process of:

- 1** Breaking the component into its functional structures, then;
- 2** Examining those parts for mere ideas, matter incidental to an idea or matters in the public domain; finally
- 3** Comparing the remaining kernel/s of potentially protectable matter with the alleged infringing program. This test is not without its critics, and the overly mechanical approach is said not to catch all infringement. In other jurisdictions, eg the UK, a less formal holistic approach is generally adopted; eg that from *Ibcos Computers v Barclays Mercantile*.

Another argument that has appeared to be offered by SCO is that it owns the IP in System V Unix, and since Linux derives (at the very least conceptually) from Unix, it is a derived work for the purposes of copyright law. SCO says that this has happened either as a result of adaptations made to its copyrighted source code, or because many programs that run on Linux also interface with elements SCO claims to own; for instance, APIs or methodologies. The former argument will be resolved on the facts in court, but the latter argument is almost too easy to defeat. Much of the supposed IP is based on general methodologies developed since the 1960s, and much that has been taught on computer science courses in universities to countless generations of students. The scope for arguing that any of this is, ever was or remains, proprietary is, I would suggest, highly limited. The suggestion that because a program uses an API or methodology, it becomes a protectable derivative work, is absurd. To the extent that there is property in the API or methodology itself, then anything other than the most trivial of implementations would probably involve sufficient work to give it a self-contained copyright that is distinct from any in the API or methodology.

However additionally much of the Linux kernel and associated code is crafted to comply with a Posix API, not some imaginary SCO API; in any event as discussed earlier it is X/Open that owns the API not SCO. As for the derived work argument, we have already seen that in reality Linux is based on the X/Open API's, public





domain code, the general state-of-the-art, BSD code and the ingenuity of Linux coders.

It is absolutely impossible to know at this stage what the outcome of this analysis would be; not least because at the time of writing, SCO still won't indicate which specific part of the (publicly available) code it claims to be the infringement.

Doctrine of pre-emption

Mark Heise, a lawyer with Boies' firm, has impeached the validity of the GPL in America in a novel way; he says *'Section 301 of the Copyright Act says the Copyright Act pre-empts any claims that are governed regarding use, distribution and copying. We believe that although the GPL is being tossed into the fray, it is pre-empted by federal copyright law.'* He goes on to say that US copyright law under section 117 of the *Copyright Act* permits the making of only one copy of software and federal law pre-empts attempts to forbid that. So what is pre-emption? It is an American legal construct governing any conflict between contracts, licences, common law, state law and federal laws (purists will argue that pre-emption applies within the EU also as between EU legislation and national law). At a very crude level, it states that where these limit the rights given by Federal law, Federal law prevails; if they try, directly or indirectly, to remove rights granted by Federal law, they are invalid. Since the GPL allows any number of copies and federal law permits one copy, the GPL is thereby unlawful. This is nonsense of the most arrant sort. The purpose of the section on any sensible reading is to give people a minimum right to a backup copy, and neither the intent of the Act nor its text prohibits greater copying. How any lawyer can offer such an argument is truly perplexing. There is, of course, the possibility that SCO has formulated some clever argument relating to the viral element of the GPL with the intent of neutering it, but if they have, they have not disclosed it nor is it easy to guess. A more realistic

hypothesis is that Heise is not a copyright lawyer. He has been given a brief and has done what research and analysis he can, but has made mistakes. Were I asked to handle an ecclesiastical case, I'm sure I'd make a proper mess of it too.

Trade secrets

I have seen it said that by publicly distributing code SCO has foregone any trade secret rights. Indeed the judgement in the BSD case, *USL v BSDi*, supports this view; in that case the judge said that header files, filenames and function names were *'...not secret since these were all available from unprotected files...'* However, the fact that code is available everywhere without a confidentiality obligation is not relevant if the overall structure of the infringed code is such that it is difficult to determine that structure, unless an existing copy is downloaded to circumvent the experimentation or development needed to get that structure from the code. Whether this argument applies to an entire operating system code is untested, but it would seem to be a viable argument in principle.

The second argument that SCO is using is that it cannot disclose which source files are infringing, since this would be tantamount to disclosing their source; at the moment, they say, it is hiding in plain view, and until they disclose it, no-one knows which bits are theirs. The difficulty with this argument is that while true only up to a point, you have to step back and ask *'what is the purpose of a trade secret?'* It is to make sure that no one can copy your source code or use it to their advantage or your disadvantage. If that code is already in the public domain, they can do this anyway; they can rifle through the Linux source taking whatever bits they want while probably being indifferent to its SCO origins. Only if an adversary were specifically interested in knowing what SCO used would this be an advantage, and it's a bit of a stretch to think up anything credible. One imagines however, that SCO would be more interested in ensuring its code was stripped from Linux rather than allowing its misuse in order to keep it secret and maintain some unclear trade secrets advantage.

“The suggestion that because a program uses an API or methodology, it becomes a protectable derivative work, is absurd.”

Furthermore, there is a principle in law that where one suffers damage, one must attempt to minimise it using whatever steps are reasonable in the circumstances; the principle of 'mitigation of damage'. Whether the continued refusal to disclose allegedly infringing files breaches this principle will depend on whether SCO's arguments that it is reasonable not to do so hold water. I'm a little sceptical given their previous statements and the general demeanour of the company.

An additional point about trade secrets is that to protect them completely there needs, broadly, to be either a relationship between the person disclosing and the person disclosed to – eg seller and buyer – or a relationship of sufficient proximity that the court would feel it equitable to

bind the recipient of the trade secret in the absence of any other relationship. That is not to say that an innocent recipient of a trade secret is automatically home free in the UK: an injunction and damages might well be had; however, in US that is the best presumption. In the US, it has been said that a trade secret once publicised is lost: ‘...once that trade secret has been released into the public domain there is no retrieving it...’ (*Religious Technology Centre v Netcom Communications*). However people have to have seen it for it to be ‘released’ which raises the issue of whether a trade secret hiding in plain sight, such as in kernel source, has been released. For historical or other reasons, some of what SCO claims to be its trade secrets have leaked into the public domain; either through general education of the



IT industry, the improper access to SCO source code or because of discovery in litigation (eg *USL v BSDi*). Where this is unlawful the remedy is likely to be an award of damages against any unlawful discloser. It is possible that a court might injunct the distribution of code but it would have to balance the effectiveness of the remedy against the effect on innocent third parties and whether it is a proportionate measure. Additionally, if the innocent kernel developers relied on the *bona fides* of ideas they received and so developed other code based on it, there is the equitable doctrine of ‘change of position’, an estoppel to defend against an injunction or damages being granted. Again, I’m somewhat sceptical that they would get anything other than damages at the very, very best.

SCO’S PROSPECTS

I’ll be honest: I don’t know if SCO will win, fully or partially. Insufficient papers have yet been disclosed via discovery to find out (discovery is the process whereby parties to a case are required to reveal the evidence they will be relying on, and pertinent material the other side wishes to see). IBM has just applied for what is a very broad level of discovery, and I imagine some of this will become available for analysis. However, the evidence is so far mixed for both sides. SCO has made some very peculiar assertions, and

failed to undertake a sufficiently clear analysis to distinguish between historic, public domain and SCO non-proprietary or proprietary code. The court atmospherics of much of this would, despite SGI’s behaviour, seem poor for SCO; though between now and the hearing, additional evidence or better arguments may compensate. Certainly SCO should not have employed a media legal star with negligible technology experience just for the PR value. The case needs a specialist technology law firm (Sorry SCO, we’re busy).

SCO vs LINUX

History of the dispute

1968

Donald Knuth writes (probably reworking earlier papers) some packet filter software. Thompson and Richie at AT&T write portions of code used in UNIX.

1985

IBM take a UNIX licence from AT&T.



1993

Novell buys UNIX source & patents from AT&T.

1994

Novell sells full UNIX Licence to Sun.



1994

Novell sells UNIX trademark and UNIX specification to the Open Group.

1995

SCO buys UNIX from Novell – some confusion over whether all patents and copyrights transferred – Novell can’t find their copy but confirm the signature is valid, so that some copyrights may be owned by SCO.

Feb 2001

SCO & Caldera merge. Later release pro-Linux statements.



June 2002

Darl McBride takes over as SCO CEO.



2002

Caldera releases code as Open source.

Jan 2003

LinuxWorld Expo: IBM makes a keynote speech about “Linux coming of age” that reportedly upsets McBride by stating their intent to “obliterate UNIX”. McBride then hires Boies’ law firm.

Jun 2003

SCO begins showing, under non-disclosure agreements, selective limited code fragments bearing a close similarity to code in the kernel.

Jul 2003

Evidence comes to light of SCO forbears directly contributing code to the Linux kernel. SCO begin offering a ‘get out jail free’ licence to Linux users. SCO continues to offer the Linux kernel from its own servers. SCO refuses to disclose alleged infringing code.

Aug 2003

Red Hat and IBM counter-sue. SCO declare their strategy will defend in part on the invalidity of the GPL, dubbed by some as the ‘Chewbacca defence’. SCO reveals two examples fragment of code which are traced to public domain code or no-copied code based on public domain specifications.

WHAT IF A UK LINUX USER IS THREATENED BY SCO?

I do not think SCO presents too much of a problem for European or UK users at present. If SCO approaches anyone in the UK, then the chances are that it will ask for a licence fee based on the argument that all (under a derivatives works

theory) or some of Linux is their copyright for which, therefore, a licence is required. In the US, Eben Moglen, Counsel for the FSF, has offered the view that one need not pay since the US Copyright Act permits the running of a

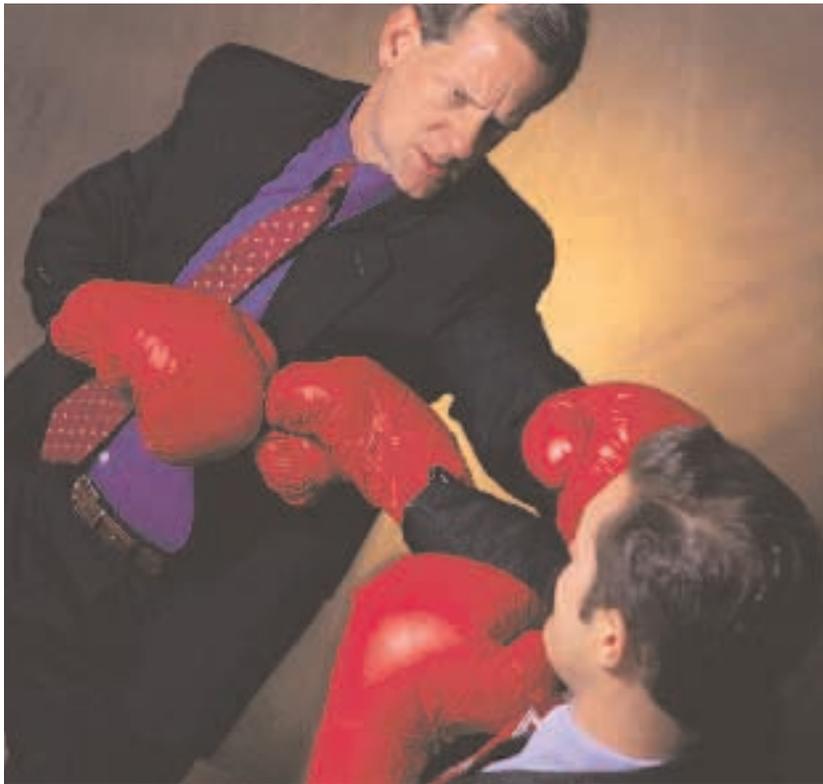


COVER FEATURE **SCO**

computer program by copying it onto a computer without the need for a licence. This may be good law in the US, but in the UK it doesn't run; section 17(6) of the *Copyright Design and Patents Act 1988* provides that copying includes transient copies such as those made by copying into RAM: the opposite position from the US. In the UK mere use of the kernel irrespective of whether the source code is used would be a copyright violation *if* SCO's contentions are right.

Nonetheless, paying up has one advantage; it will at least alleviate you from the risk of litigation, but the price is a heavy one since SCO are requesting \$700 for single-user licences. Multiply this across an enterprise, and it will make Windows look cheap. Nor does SCO promise you your money back if the licence turns out to be worthless, indeed it explicitly says the opposite. They exhort users to act honestly and with integrity without feeling the need to do the same: if they were as confident as they claim they would put their money where their mouth is. Of course there are arguments to make that in these circumstances any contractual licence is void. Under the doctrine of frustration for example one might argue that the purpose of the contract was frustrated because the alleged SCO IP rights forming the basis of the contract had vanished and that SCO therefore has to reimburse the contract price.

Were I to be approached to buy a licence, I would ask them to provide detailed proof of the alleged infringement including a listing of infringing portions of the code; tell them to sue me; or wait until an English court declares that there is an infringement. If threatened with litigation, I would also seek discovery of the alleged infringing code with a view to validating the claim of infringement or getting that code removed to provide a 'clean' kernel. This is of course at the heart of the dispute between the FOSS community and SCO, and it is this that is the proof that the whole SCO circus is a protection racket. An honest company, on discovering their code had been stolen as SCO alleges, would be keen to sue for damages and an injunction; to protect their IP they would also be only too anxious to get it removed and for infringing



products withdrawn and destroyed. This is all standard fare for IP actions. SCO however has proven that it is a shyster by attempting to maintain the supposed infringement in the hope of extracting ongoing licence fees from a wide variety of sources. If they sued for damages and won they would probably only get a one off damages payment of X million dollars; perhaps the calculation is that they can get much more by extracting licence fees from all current and future users. This Machiavellian goal would, however, be shot to pieces by FOSS hackers doing the honourable thing and removing any supposedly infringing code; indeed an honest victim of copyright infringement would insist on it.

COUNTER LITIGATION

Linux is of course an international phenomenon, and the SCO case is a US one; but the substance of the case is not such that it that would prevent SCO repeating litigation in Europe or elsewhere. The result may not be entirely the same, though an analysis is beyond the scope of this article. LinuxTag in Germany has already obtained an injunction against SCO Germany for unfair competition, since SCO is engaging in FUD and not disclosing proof of the alleged infringement. SCO's failure to respond substantially so far is consistent more with its focus on litigation in America, and containing the issue elsewhere. SCO has a litigation cash pile of about US\$10M dollars and it would be nice to see that eaten up by Linux users around the World taking SCO to court; though one suspects Microsoft would merely find an excuse to buy another US\$10M licence. As an intermediate step this litigation in Germany is a good thing and it would be nice if it were happening elsewhere; I'd be happy to provide *pro bono* advice to any company wanting to do this in the UK.



Criminal sanctions

I have seen comments from some people who (understandably, but a little unrealistically) want to press criminal sanctions against SCO. Putting the share-pumping allegation to one side and considering just the issue of allegations of copyright infringement, I'd regard this as wildly unrealistic. While many would like to see Darl McBride share an intimate, if coerced, moment with Bubba and his friends in the showers of a Utah penitentiary, it is unlikely to happen. Suggestions have been made of blackmail, fraud, extortion and such like. These certainly capture the feelings of Linux users, and at a colloquial and pejorative level they'd be accurate. Unfortunately, the standard of proof required to prove criminal allegations is quite substantial: *beyond all reasonable doubt*; and I doubt if sufficient proof exists to meet that burden. It would be necessary to demonstrate that what was said was false, known to be false (or without any belief in its truth) with a dishonest intent to deceive and obtain a financial advantage. No, I do not think so, unfortunately.

CONTRIBUTING CODE TO FREE SOFTWARE

The risks

Finally, we come to honest participants in the FOSS process. SCO started off – albeit reluctantly and through necessity – in this process several years ago. It was not until Darl McBride joined SCO and tried to hijack Linux JFS and NUMA that significant thoughts about the resilience of the process arose.

Part of the problem here is to avoid adverse consequences, such as SGI's possible malfeasance in relation to the inclusion in FOSS of other peoples code by stripping out copyright notices. The issue is how to avoid such things in future. In part Google, education and open mailing lists are the answer. Linus Torvalds has said he is well aware of the need for clean code and accountability, and I have heard that the process for inclusion of code in the kernel involves elaborate and lengthy public discussion along with assurances of code provenance. This is clearly a good – if minimal – approach. If someone does include code that they should not in a FOSS project, the project team may all be liable for copyright infringement in certain circumstances. It may seem a bit bizarre to say that all x-thousand kernel developers may be liable if copyright infringing material appears. In principle however, there is an argument for finding joint and several liability attaching to some or all involved members of a guilty project. Of course, for such a large project as a kernel, this would seem an unreasonable and severe approach, given the massive size and complexity. A more realistic approach would be for liability to attach to all or some of the members of a

subsystem team, eg the USB or Ext3 team. Scary? Yes somewhat. Education has a vital role. If the SCO affair has a positive function for Linux, it is to reinforce the message that not everyone likes FOSS. Some see it as a lethal and pervasive challenge to their business that must be killed by all means fair or foul. Well then, it is understood.



ABOUT THE AUTHOR

DAVID HARRIS IS AN IP barrister, practising at www.ukitlaw.com. When not suing people, he writes 'depressingly poor code'.

THE CONTRIBUTION PROCESS

A collection of coders' caveats

For any collaborative project of mine, I would do or attempt some or all of the following (though this is a minimal approach);

- Create a rigorous but minimally complex process for accepting contributions: knowing exactly who one's contributors are and who employs them. Essentially, this would be for the purpose of flagging up potential issues such as confidentiality, competence and ownership of code.
- Undertake an educational process for private contributors about the importance of copyright and other IP issues. This would be public, and all contributors would be invited to view it. While far from a defence, it does at least set the correct parameters.
- For corporate contributors with legal departments, I would engage them in a proper process of auditing submitted code. SGI did this with XFS for example, and it delayed its contribution for quite a while; likewise IBM with its kernel contributions. In both these

cases I have faith SCO will not prevail for these reasons.

- My own preference for code that has no potential to be commercialised would be to assign copyright to the FSF. This does not, under English and US copyright law, require formalities beyond putting down in writing a clear intent to transfer fully permanently and irrevocably all rights to the FSF. Then the copyright owner must physically sign the document and should send it snail mail to them together with a disk of the code.

For code where there is, however remote, the potential to commercialise, I would not assign the totality of copyright. The legal options are numerous: granting joint ownership, offering the FSF a contractual GPL licence, creating a commercial derivative and assigning the Free version to the FSF plus a number of other variants.

- If creating code based on the works of others, I would attempt to use publicly available

non-confidential documents in order to obviate trade secret accusations: Usenet postings, mailing lists would be good for this purpose.

- When coding to standards, specifications or APIs, I would attempt to evaluate whether the owner claims them to be nonetheless proprietary. Obviously there is no generic approach to this, and the options would depend on the circumstance and who the owner was. People such as the ISC are probably OK, but asking Microsoft for formal consent to follow its 'standards' is going to be rather problematic; however that has not yet hampered *Samba* or *OpenOffice.org* (though that may well change). Whatever the response I would print off and archive the response as I would to any issue that may require proof several years later.
- Pray to St Isidore (believed by some to be patron saint of computing and the Internet) and a broad selection of the most powerful gods I can find that there are no software patents. ■

JBoss takes on Europe

When it comes to Java application servers, surely everyone in the Open Source world has heard of JBoss. But it isn't merely a success to 'free software' aficionados - the JBoss Group, a services company based around professional support of the software, already has customers including Motorola, AMD, Wells Fargo, Hitachi, Siemens. And Playboy.

Following the announcement that the group has opened offices in the UK, Switzerland, France, Finland, Germany and Ukraine, we got to speak to Sacha Labourey, the new European General Manager

LXP: Why has JBoss decided that it now needs such a major presence in Europe?

SACHA LABOUREY: The goal is clearly a commercial one. We wanted to be closer to both our customers and partners, and if needed to tune our offerings for the European market. Also it makes some things easier, even simple things like being able to make a contract with a European company in Euros.

Also, many people didn't know of the JBoss group in Europe. The software is fairly well known, but the link to the JBoss group wasn't.

LXP: You mention the potential to create a European 'flavour' of your services. Do you think there is a difference in what European customers expect?

SL: Not necessarily. What people expect is production support and to be sure that the solutions are robust. These expectations don't really change.

LXP: And there's no difference in the services you are offering.

SL: Not at the moment, though this may change in the future. One of the strongest points of our solutions is to provide high quality support, because all the core developers are employed by the JBoss group. The best argument we can give in favour of our production support is that the people who implemented the features will help the customers with problems. This support is distributed

NICK VEITCH investigates the JBoss group's decision to open offices all over Europe.

amongst all the developers, whether they are based in the US or in Europe.

Many of the developers are live and work in Europe. There seems to be a common perception that JBoss is very American, but a lot of the development is done in Europe.

LXP: Obviously JBoss grew up out of Open Source and follows that model of development. Linux Pro supposes in common with other organisations that have built services on top of 'free' software that the continuation of JBoss as Open Source is important to you from both a business and ideological standpoint?

SL: You are right. We are an entirely services-based company. We call it professional Open Source. The goal for us is to hire the main developers, specifically in Java application servers and the surrounding technologies. This explains the hiring Gavin King, formerly of Hibernate, and Remy Maucherat from Tomcat, for instance. That's really what we want to do. Instead of people developing as a hobby and not be able to give you real help when you need it, we want to professionalize that. Perhaps the difference between what the JBoss group and other projects is the

The JBoss website at www.jboss.org has had more downloads than Sun's J2E.



model we have chosen – there are all kinds of models, for example, projects funded by government or universities – the goal there is often not to run a business or provide a service. Obviously there are companies like MySQL who provide services and licences, because the software is dual licensed. JBoss doesn't sell any licenses or subscriptions – just services. In that case our clear offer is production and development support, and training.

Some people ask if JBoss is a viable business. We reply with the fact that we have no external funding, this is a self-grown business, which is profitable.

LXP: We at Linux Pro think the myth that you can't build businesses around Open Source software died several years ago. The JBoss software certainly seems to be incredibly popular.

SL: That's correct. JBoss has had millions of downloads. We consider ourselves to be the reference implementation. Sun's J2E has had 2 million downloads, but we've already gone significantly past that.

Users download it to try the spec. In Europe, one major noticeable difference to the USA market is that there are lots of public sector clients looking at Open Source. We also see more and more finance companies using JBoss for testing. They may use Websphere or whatever for their main production, but they often experiment with JBoss for stuff in development.

While a few years ago JBoss was really only used in testing environments, we see a change now. Customers that are not obviously Open Source fanatics at the start are now much more keen, and that interest is infectious, leading to their exploration of other Open Source solutions.

LXP: Do you have a feel for how many people choose to implement JBoss specifically on Linux?

SL: It's very difficult to tell, as downloads and usage numbers can be disparate. Our research tells us that it is probably about a third, with Windows at about a third and everything else at a similar size.

LXP: Sun clearly seems to think that outside of Solaris, the Linux version of Java is the most solid. With Microsoft having some difficulties with Java, it may be interesting how that plays out.

SL: The Linux JBoss couple makes a lot of sense. I would say that the JVM on Linux isn't at quite the same level of quality as Solaris yet. It isn't really a problem with Linux or Java, but the mix. You have to fine-tune a lot to find the best kernel, libraries and so on. It takes more tweaking.

When this changes, it will be a tremendous boost for Linux and JBoss. I think it will change because Sun is keen on Linux now; also the distributions are improving the JVM support, and now there is a new threading model in the kernel. All these things should make a difference.

The Linux and JBoss combination is an obvious choice. If you want to reduce costs and use a full Open Source stack, there aren't many other options.

LXP: Another thing we wanted to ask about is the JBoss 'boot camps'. Presumably they make sense for you...

SL: Clearly it's a way to brainwash people to use JBoss! No, seriously, we've found that it's a good way of effectively evangelising JBoss and answering people's questions. We get a lot of IT directors from big companies who come to see what all the fuss is about.

We also get people that already use JBoss who just want to talk to the developers about the future. There is no real single category, no single type of attendee. We recently ran our first boot camp in Europe and we will certainly be doing more.

LXP: I wanted to ask about the JBoss IDE as well and the future plans for that. There seems to be a fair amount of competition in that area.

Is it still worth developing that or wouldn't it be better if everyone just settled on one open source tool and worked together on that?

SL: Well, since JBoss is essentially a rebranding of Eclipse, we are not trying to re-invent a full IDE because it would be a waste of effort. There is Eclipse, Netbeans, and JBuilder already. The goal was to have a better level of integration of JBoss with Eclipse. With JBoss 4.0 we'll have the AOP features. Inside the source code you can add the tags that say 'this is a webservice', 'this must be cached' and so on. Then you don't need deployment descriptors. For this a high degree of integration is required.

“The Linux and JBoss combination is an obvious choice if you want to reduce costs and use a full Open Source stack.”

LXP: When is JBoss 4.0 due for release?

SL: I don't have the roadmap in my head, but towards the end of this year. That will really bring some excellent new tools to various types of developers. There is a huge category of developers who don't care whether things meet the J2E standards or not, they just want to quickly solve problems, and this will be great for them.

LXP: The only negative thing I have heard about in relation to JBoss is that the documentation is weak.

SL: We have two kinds of documentation – the free documentation and the paid-for documentation. The latter is very deep, so you will find what you need, but it isn't as friendly for beginners. There is a companion book by O'Reilly which I wrote with Bill Burke, which has everything you need to start programming.

The official free documentation – we have tried to do something to improve it in the past, and I admit it was bad first time around. We are presently in the process of totally rewriting it from the ground up. We have someone in the UK working on that now.

LXP: We look forward to seeing that. Is there anything else you'd like to add?

SL: The key thing for us is that the JBoss group wants to promote the idea of professional Open Source throughout the industry, and we want to give the best support we can. We take it very seriously. ■

Monkey see, monkey do

Novell's revolutionary plans for enterprise Linux

LINUX PRO: How would you say Novell's attitude to Linux has changed in the last year?

BRIAN GREEN: The first thing to point out is that Novell has been delivering solutions that sit on top of Linux for some time now. In late 1999 Novell announced its vision of "One Net", which is making sure that customers have access to information anywhere, any time, any place, and even on any device. On the back of that, we also said that we would enable that with the infrastructures that our customers have. By implication, that meant that Novell had to embrace our customers' platforms, whether that be Netware, which is obviously something Novell is well known for, but also Windows, Solaris, and Linux.

To enable us to achieve that, we ported some of our technologies to Linux to make that happen, and we've been doing that for some time. The most well-known one was our directory service *eDirectory*. However, the reason things have accelerated and we've become a lot more vocal is because our customers have asked us to. We have customer councils around the world where we take feedback from our customers and potential customers about what's driving them in their organisations and what kind of technologies they are looking for. Late last year it became more and more apparent that we had a new, growing theme that Novell had an opportunity to address and we felt we had value to add – that being the Linux platform.

LXP: So you'd say it was your customers that attracted Novell to the open-source movement?

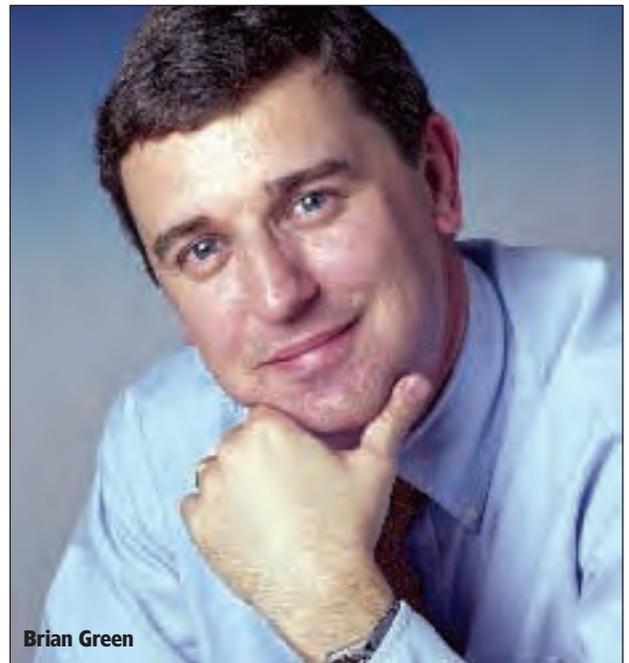
BG: Yes, absolutely. You know, we're a customer-driven company – we produce solutions to meet our customers' demands rather than because it's vogue. The reason we're doing it as an organisation, and the reason we're accelerating that is because of our customers' demand. In addition, some of our customers came to us and said "We like the opportunity of Linux, the cost reductions that Linux can bring to our organisation, the openness and choice it gives us – we're not tying ourselves to any particular vendor – but there are some inhibitors we wish could be addressed." So, Novell saw an opportunity for us to help our customers' organisations remove those inhibitors, and that's really what we've been doing with our change in attitude around Linux – looking at the inhibitors our customers have and seeing what we can do to remove those inhibitors.

LXP: Where does Ximian fit into these inhibitors?

BG: Well, Ximian has a similar attitude to Novell – its vision was to enable Linux in the enterprise. Our customers are asking us to enable Linux in the enterprise, and Ximian has technology that complements Novell's, and also enables us to accelerate solution delivery. For example, Novell has a

Having acquired Ximian, launched a Certified Linux Engineer program, and announced it intends to port its services to Linux, Novell is clearly redefining its role in the enterprise Linux marketplace. **PAUL HUDSON** spoke to Brian Green, the Director of Enterprise Solutions for Novell EMEA, and hears no evil...

Novell



Brian Green

product which is desktop management for Windows, called *ZENworks*, that we've been very successful with. Ximian has a product called *Red Carpet*, which I'm sure *Linux Pro* readers are familiar with. The opportunity exists to bring these two solutions together to offer desktop management that's policy-based, based upon an individual's directory entity irrespective of whether they are running Windows or Linux – that's something of great interest to our customers and Ximian helps us to achieve that faster.

LXP: So if Novell are integrating Ximian's products with Novell, will the Ximian product line be discontinued?

BG: No, no, definitely not – we are committed to take the Ximian technology lines forward, but what we *will* do is to integrate Ximian technology and *ZENWorks* technology together, so that those organisations that want a management solution that embraces both Linux and Windows can seamlessly manage that from a single point of administration. So your identity is common whether you're sitting at Windows or Linux. Make no mistake, Novell didn't acquire Ximian to terminate it – we bought it because it's a fit for our organisation, and it helps us meet the requirements of our potential new customers.

LXP: When you bought Ximian, you also picked up *Mono* – where does that fit into your plans?

BG: If you look at the *Mono* project, one of the things that *Mono* will help customers do is to increase the applications

that sit on the Linux platform. One of the inhibitors that some organisations have come back to Novell with is the lack of applications; although it's true to say that the number of ISVs supporting Linux today is growing. One of the things Novell wants to do is help accelerate that, and obviously the ability of the *Mono* project is to potentially take advantage of Microsoft's development tools and development framework so you can develop on Windows and deploy on Linux – this opens up a number of opportunities for developers in a corporate environment, and also offers an opportunity for Novell, through the *Mono* project, to talk to ISVs and assist them in porting apps to Linux. It's a project that's currently Open Source, and Novell is continuing to invest in it. Going forward, there will be solutions based around that project.

There is one question that Novell has been asked quite a lot since we acquired Ximian, and that's about the *Mono*.NET patent issue. The Microsoft .NET framework technology is based upon the patent technology that comes under the ECMA/ISO submission by Jim Miller at Microsoft, and Jim Miller is the patent owner as a Microsoft employee. Jim Miller has made some comments to the fact that it is absolutely their intention that C# and the CLI will be available on a royalty-free and otherwise reasonable and non-discriminatory (RAND) basis.

LXP: What brought about the introduction of the Certified Engineer program?

BG: The Certified Linux Engineer program came from Novell addressing two needs. This is not the only thing we're doing in this space, but when we ask our current customer-base about Linux, we always get the same two questions about skillsets. The first is "How can you help me become more skilled in Linux?" – and Novell is addressing that in a number of ways: through self-study training materials that we're giving away for free, hands-on technical training around early-adopter technology that we'll introduce early next month; and our Certified Linux Engineer program.

Also, the second question is "We recognise that Linux is not new, and there's a numerous amount of linux professionals available on the market, but how do I measure a Linux professional? And if I want to hire someone that has a Linux background, how do I know that their skillset is relevant for my enterprise?" So one of the things that the certification brings to the Linux community is an industry-recognised certification backed by a company such as Novell, which arguably invented IT certification, so that an IT director can hire someone with confidence that this individual is able to perform the task that they want.

So really the Certified Linux Engineer program is about the need for developers to demonstrate his or her skillset, and also to increase the skillset of IT professionals already in the company around the new Linux platform.

LXP: How committed is Novell to giving code changes and improvements to the community under the GPL?

BG: Novell will stand by its reputation. Even prior to our acquisition of Ximian, Novell was taking technologies we had inside our organisation and making them available under the GPL, such as a UDDI server that sits on the server for

identity discovery, and various technology that manages farms of *Apache* servers. Obviously, we continue to support the GNOME project through Ximian, and also *Mono* – these projects continue to be Open Source.

LXP: To what extent is Novell porting its services across to Linux?

BG: We've issued a lot of announcements in the last few months, but the one that probably drew the most attention was that in April 2003 we said we'd even take *Netware*, and all the services you see on there today, and port that onto a Linux distro. Since then we've also made announcements around our collaboration product, *Groupwise*, so that will also run on a Linux distro. We announced that at LinuxWorld San Francisco in late August – around the time we made our acquisition of Ximian. As we integrate our *ZENWorks* with *Ximian Red Carpet*, obviously we will continue to support that too. We also have web services technology and identity management technologies. So, yes, in a nutshell, we're taking technologies we've been producing within our company on multiple platforms and making sure that they embrace the Linux community and the Linux platform to meet the requirements of our customers.

LXP: How much growth do you anticipate in the Linux market over the next few years?

BG: It's true to say that the Linux market is growing, both on the server and on the desktop, but the most important thing to use is that our customers are saying that they are embracing Linux and moving to Linux. Novell is an organisation that wishes to make sure we meet the needs of our customers, and, of course, it's also true that we see this an opportunity for increased revenue for us as an organisation while we're meeting the needs of our customers.

Linux is a growing operating system, so we need to address it. That doesn't mean to say we're moving away from the other platforms we support as a company as well,

“Even prior to its acquisition of Ximian, Novell was taking technologies we had and making them available under the GPL”

however it is true to say that Linux is the fastest-growing operating system platform on the market today, and we see that as continuing.

LXP: What are Novell's Linux plans for the future?

BG: We've made aggressive roadmaps to port our network and services that you can see today, such as the *Netware* kernel that currently runs on Windows, so as we're moving to the next calendar year you'll see things like *Groupwise* on Linux shipping in the first half of next year. You'll also see our web services and identity management solutions – more of those solutions will become available on a native Linux environment. If you go back to the collaboration environment you'll see Novell being able to offer a collaboration environment that runs completely on Linux from both a server perspective and client perspective. ■



HOSTING

WEB HOST FOCUS

Host Europe

Back in the mid 1990s when many SMEs were discovering the power of a website, many companies chose Magic Moments – the original part of what has since become one of the companies making up Host Europe. Back then, Abby Hardoon (now CEO) often handled the tech support questions himself. Now in 2003, Host Europe is a much larger company made up of Magic Moments, Webfusion, 123-reg, and Dedicated Servers, and these subsidiaries provide a range of web hosting solutions designed to cater for all needs.

Hosting solutions start at the most basic level with 123-reg, where, for the equivalent of just £1.59 a month you can get 100MB of web space, 15 email boxes, and a 1.5GB monthly transfer limit. Moving up the chain, the the WebFusion Home account costs £9.95 a month and includes 1.5GB of web

PAUL HUDSON investigates one of the industry's longest-running providers in the third part of our series examining the most popular web hosts...

space, a 10GB monthly transfer limit, and has MySQL/PHP support for people who want to add dynamic pages to their site. For resellers looking for an all-in-one solution, both WebFusion and Magic Moments offer reseller accounts for £50/month, which allow you to host any number of domains, with a 10GB monthly transfer limit per domain.

Serving corporate users and other sites that have high traffic, Dedicated Servers' range starts at just under £100 for a basic Linux server. For that princely sum you'll get a 2GHz Celeron with 128MB of RAM and a 40GB hard drive, as well as 30GB of data transfer every month. From that £100 upwards there are a variety of solutions offered, with the top-end being the Linux Corporate Advanced solution, which, for £6000 a month, gets you four web servers, a mail server, and a database server, as well as a firewall, a

MAGIC MOMENTS

Linux Pro spoke to Abby Hardoon, the CEO of Host Europe, about what makes the company magic...

LINUX PRO: Tell us your unique advantages over your competitors?

ABBY HARDOON: We've prided ourselves on our innovative approach to web hosting since the very beginning. We were the first company to develop a flexible, in-house hosting control panel, which as well as proving extremely popular with our customers, gives us the ability to make specific improvements based on customer needs. *MyServerWorld10.0™* was recently launched with great success, and marks a huge step forward in giving customers total and flexible control of their shared hosting solutions. We fully intend to continue this trend of innovation within the hosting market, both in terms of the products we offer and the service we deliver.

Our experienced Sales team is specially trained to work with our customers to ensure they find the hosting solution that best matches their needs, technical resources and future intentions.



“Linux has already played a huge part in Host Europe’s success... and strong Linux-based expertise of our tech staff allows us to support a broad range of servers.”

ABBY HARDOON, CEO HOST EUROPE

Our product range and technical expertise is broad enough that customers need never leave us. We have long-term customers who originally simply registered a domain name with us and whose businesses we have been delighted to watch

grow through our hosting products up to high-end clustered server packages over the years.

LXP: What would you say is your key target market at present?

load balancer, and 1200GB of data transfer each month – an awful lot of kit for what is quite a reasonable price tag.

Clients and computers

It's clear that Host Europe hopes to offer something to fit your exact needs, whatever they are. This is underlined quite strongly by their high customer numbers – they host just under 600,000 domain names for over 130,000 customers worldwide. In addition, Host Europe has about 4,200 dedicated and co-located servers spread across its data centres. The London data centre alone currently has 190Mb/s bandwidth available, delivered by three separate providers, of which, surprisingly enough, only about 40% is being used at any one time on average – this means that, even under more than double the average traffic, customers will still get their full share of bandwidth – something that few other companies can promise.

To back up their hardware performance, customers are also given the choice of support options to along with their server. Dedicated Servers, for example, has a comprehensive SLA promising four key things: your server will be rebooted within thirty minutes of you requesting it, your server will be available online a minimum of 99.9% of the time, your price will never increase even if the same service is more expensive for new customers, and also that every customer will always have their full amount of bandwidth available. They also get 24x7x365 technical support by both telephone and email, which means help is never far away if you find yourself stuck.

For WebFusion, Magic Moments, and 123-reg, there are other service agreements, although you always have the choice of picking the service level fits you best. All Linux servers come with Red Hat Linux pre-installed and ready, and also includes an advanced web-based control panel where you can configure email addresses, customise your firewall, purchase upgrades, and more – you may even find you can solve your problems yourself and pocket the cash! ■

DEDICATED CUSTOMERS?

James Puddicombe, MD of JJP Solutions, expresses his thoughts on Host Europe...

An e-commerce sales channel for a business has to make more money than it costs to run. It is blindingly obvious that a Linux platform provides more power and reliability than Windows for the same budget, so when we launched JJP Solutions as a web design and site management company working on a commission-only basis, we made sure our systems were based on Linux, in fact on LAMP.

But the second stage of the problem was finding a professional, cost-effective hosting provider. We experimented with some in the US – initially attracted by the lower bandwidth costs than the UK. The higher latency was bearable, but the professionalism of the hosting companies we tried was non-existent, from not realising the server hardware was faulty, through to having their status page in the same location as the rest of their servers.

Fortunately we had retained our foothold in the UK with Host Europe. We've become ever more impressed with their dedication to customer support. They have an empathy with our problems so that when we consider expanding even a simple hosting architecture their team don't just try to flog us boxes but instead end up working alongside us to draw-up the platform blueprint.



Our business has taken off – swimming against the tide of conventional e-commerce services. In fact in the last 24 months, we've grown from three virtual servers through to a cluster of four dedicated machines providing local real-time backup and fully automated off-site disaster recovery facilities – all with the help and support of the Host Europe team. From our experience over the last 12 months, it's obvious that commission-only e-commerce services will become the standard for thousands of companies around the world. Open Source solutions and the Linux community, working alongside professional companies such as Host Europe, will be key to the undoubted success of this proposition.

AH: Host Europe caters for a broad range of customers, from those creating their first online presence to enterprises who demand cluster solutions and powerful firewall and monitoring solutions. A key part of our customer base is the SME (small- and medium-sized enterprise) market, and our product range is designed so that a company can register a domain name and establish a simple website as they start their business with limited resources.

Gradually, as businesses become increasingly successful, customers can upgrade to more flexible hosting including, say, an e-commerce package from WebFusion, and eventually move up through the range of dedicated servers to powerful multi-server solutions.

LXP: What kind of hardware do you deploy?

AH: Host Europe's data centres are home to thousands of servers across a broad range of

hardware, including Sun, IBM, Compaq, Dell and HP. We also have a close relationship with Sun Cobalt, hosting more than anyone else in Europe of the popular and flexible Cobalt RaQ family. We use powerful Cisco 7206VXR routers and award-winning Enterprise Alpine 3808 Ethernet Service Provisioning Switches to help manage our network. We are committed to our policy of sourcing the best quality hardware to guarantee our customers the best and most reliable service possible.

LXP: How much does Linux figure into your hosting plans?

AH: Linux has already played a huge part in Host Europe's success, and we fully expect this trend to continue. From the company's inception in 1995, all of Host Europe's shared hosting has been Linux-based, and the strong Linux-based expertise of our technical staff also allows us to fully support a broad range of Linux-platform dedicated servers.

LXP: Where do you see the hosting market being in three years time?

AH: Over the last couple of years the hosting market has been dominated by mergers and acquisitions, and it seems likely that this trend will continue. We feel that as the market consolidates and the customer rightly demands more and more of the provider, the value-added aspects of a hosting provider's service will become more important.

For example, at Host Europe we place a strong emphasis on customer service, and have recently implemented staff and procedure review and training processes to ensure that the service our customers receive remains second to none. We have recently added spam filters to our popular Linux-based WebFusion and Magic Moments packages, and developed a comprehensive new FAQ section for 123-reg and a support site for customers of our other brands, enabling clients rapid access to the information they need.