

LINUX

THE UK'S BEST-SELLING LINUX MAGAZINE **FORMAT**



KERNEL 2.6 IS BORN!

All the new features, improvements and tweaks coming to a Linux box near you – the next generation of computing unleashed **p43**

HOW TO:

- Fast installing websites **p46**
- Use regular expressions **p49**
- Create parallel threads in Bash **p78**
- Rebuild a disk using LDF **p82**
- Set up an Apache 2 server **p88**

GROWING THE CROWD

Why a Linux world may be a good thing to have **p98**

THE BEST JAVA IDEs

Choosing what's best for your development tool chest **p103**

HOT HARDWARE

Amarr, Golden Wonder, Wonder and Flight is how 2000 were reviewed **p117**

WHY YOU WILL CARE WHAT'S IN PHP 5

2008 PHP 5.0 and more reviewed by Jeff Stupp **p118**

ZOWIEWINDER

Can't find a link? Use a Linux web browser to do the job **p120**



Our new photo recognises up to 50% of Windows! Be first to know on USB Linux **p123**




LXF48 NOVEMBER 2009 £4.99
 www.linuxformat.com

WWW.LINUXFORMAT.CO.UK

Quality control

Being a Linux magazine, we tend to face a rather different set of, for want of a better word, 'challenges', to other magazines. If this were *The Official Windows XP magazine* (well within spitting distance from where the LXF team sits), we would know months in advance the exact date when various major software releases were due, and by and large, those dates would be accurate.

In the world of Linux though, you can throw that kind of certainty out the window. With the feature freeze on, and the kernel already up to test version 4, we have decided now is an appropriate time to tell you about what it will actually mean for the average Linux user. But of course, we still don't know for sure exactly when it will be finished. There's only one man on the planet who could know, and he probably hasn't even made up his mind yet. The simple fact of the matter is that most Open Source software is released when it's finished, not when some marketing department reckons would be a good date. Far from being a disadvantage (except to

unfortunate magazine journalists), this should be seen as one of the most positive strengths of Free software development.

Also on the 'quite not there yet' list is PHP5, but now is definitely the time to start getting excited about it. If you are embarking on a large web-based project, you may want to hear what Zeev Suraski has exclusively revealed to us about the upcoming release.

But Linux is as much about the here and now: hardware-wise, as well as our new Linux-compatible peripheral review series, we have an Opteron-powered workstation and a dual-Xeon server on test; and keen coders will want to take a look at our *Roundup* comparing the best Java development environments.

If your tastes lead you more to the practical implementation of Linux, check out our *Server School* tutorial for details of setting up Apache 2, or why not try your hand at dynamic particle systems in *Blender*? Whatever your interest in Linux, we hope you enjoy this issue!



Nick Veitch EDITOR

Find out what makes the new kernel ready for next-generation computing **p43**

A new logo, a new slogan and a new outlook – we take a detailed look at the new SUSE **p52**

Everything you need to know about the next version of PHP **p56**



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
When not helping out novices in the mag and on the forum, Andy still finds plenty time to write the news!



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



Richard Drummond
At one point we were taking bets on whether the kernel would arrive before Rich's text...



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



Paul Hudson
In spite of the burden of being lead architect of SKYLang, Paul still has time to write other stuff for the mag. Eventually.

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

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

Neil Bothwick
When he's not sourcing the very best software for our discs, we don't know what he gets up to – nor do we care!

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

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

CONTACT US

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

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

LXF46 November 2003

Contents



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



COVER HIGHLIGHTS

The new Linux kernel 43

What's in it, why it's happening, and what it means to you

17 Hardware reviews

A workstation, a server and Linux-compatible keyboards

32 ROUNDUP: The best Java IDEs

The top development tools on test!

50 DEVELOPER PROFILE: ZoneMinder

Philip Coombes celebrates a successful prosecution!

56 PHP 5 – EXCLUSIVE PREVIEW!

PHP guru Zeev Suraski spills the beans



“Open Source is a way of seeing the world... perhaps even after spending **all that money** promoting ‘Linux’, **IBM might not ‘get’ it.**”

09 Hoyt Duff doesn't believe the TV hype!

‘The performance enhancements in 2.6 should benefit all users, no matter the role in which we deploy Linux...’

43 The new kernel puts the ‘k’ in ‘kool’...

LINUXPRO

Linux Expo UK 2003 Show Guide! See you there, folks! **Red Hat Certification** Distro-specific trainin **Beat the Hackers!** 7safe shares its expertise with you **Servers & Storage** Xinit leads the way...

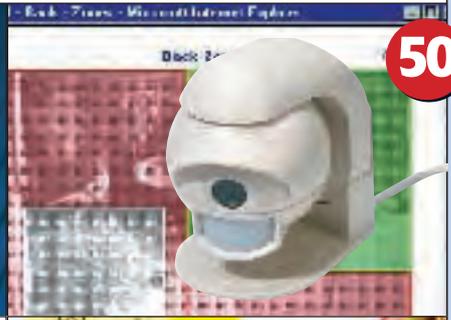




18 The Pro3D looks as good as it performs!



<p>6 News Analysis of the latest Linux stories</p> <p>14 Mailserver Your points of view and tips</p> <p>18 Armari Pro3D review 64-bit workstation stunner</p> <p>20 Fujitsu Primergy TX300 server Stealthy with good performance/price</p> <p>22 OpenOffice.org 1.1 review Feature-junkie's wish-list comes true!</p> <p>24 Aethera review Open Source work collaboration suite</p> <p>26 Linux peripherals NEW SERIES Compatible hardware</p> <p>28 Book reviews Essential treeware for Tuxters</p> <p>38 HotPicks The best new Open Source software</p> <p>35 EXCLUSIVE: The new-look SUSE The Linux Experts: Simply Change</p>	<p>56 EXCLUSIVE: PHP 5 A look ahead with the CTO of Zend</p> <p>58 What on Earth is "The GRID"? The future of computing...</p> <p>64 Beginners' tutorial Spoofing IE-only websites</p> <p>68 Regular Expressions tutorial Their importance is fundamental</p> <p>74 Compiler writing – flex & bison Handling conditional statements</p> <p>78 Blender tutorial Smokin! Render some Particle effects</p> <p>82 The GIMP tutorial Restore your heirloom photographs</p> <p>84 PHP tutorial Output buffering saves bandwidth</p> <p>94 Answers Your Linux queries answered</p> <p>110 User Groups Local support and fun</p>
---	--



COVERDISCS

A DVD or 2 CDs packed full of the latest Linux goodies **104**



CDS A AND B

Gentoo Linux Great distro packs a big punch; **Kernel** latest version at time of printing; **Nano** The Emacs or Vi argument is over; **Mass Rotate** View landscape and portrait digital photos on your TV; **Seapine Surround SCM FULL** source code system worth £395; **BZFlag** Multiplayer 3D tank battle; **Jahshaka** Real-time multimedia editor; **Kroupware** Open Source MS Exchange replacement



DVD

Gentoo Linux Extra packages and libraries; **RUNT** Complete network testing distro; **RPM Analyzer** end your dependency problems; **VCool** chill out overheating chips **and much more!**



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

SAVE MONEY!

SUBSCRIBE TO LINUX FORMAT



See page 102 or phone 0870 4448645

Newsdesk

- IBM Linux TV ad ● New GNOME releases ● Red Flag Linux ● XFree86 forked ● SUN's Mad Hatter
- Desktop Linux Consortium ● Telstra broadens Linux test ● New YOPY Linux PDA ● Petition the RIAA

SHARED ACCESS

Samba gets active

Jeremy Allison's *Samba* team has entered the final phase of development in readiness for the release of *Samba* 3.0. The first release candidate hit the mirrors early in August, but no firm date has been set for the final release at time of writing.

Samba, which allows Linux and UNIX users to access and share data with Windows computers using the Server Message Block (SMB) protocol, is regarded as one of the key technologies driving Linux forward in corporate circles.

The third major release significantly adds support for Microsoft's Active Directory, allowing users to join an Active Directory Services realm as a member server and authenticate users using LDAP/Kerberos. There is also a new internal authentication system, improved support for printing under Windows 2000, XP and 2003 and tools for migrating from a Windows NT 4.0 domain to a *Samba* domain while retaining user, group and domain Security Identifiers (SIDs).

Early tests suggest that the new version brings significant performance gains, with IT Week claiming the server can handle "four times as many clients as a similarly configured system running Windows 2000 before performance begins to drop."

The GPLed software is also being pushed centre stage in the long-running SCO vs IBM case as, at a recent SCO event, CEO Darl McBride highlighted *Samba*'s inclusion in the company's latest *UnixWare* products, even though SCO lawyers had previously declared the GPL invalid,



Jeremy Allison: "For SCO to continue to use Open Source/Free Software while attacking others for using it is the epitome of hypocrisy."

and McBride himself had said the license wasn't about making software free but "destroying value".

In an open letter, the *Samba* team wrote: "We observe that SCO is both attacking the GPL on the one hand and benefiting from the GPL on the other hand. They can't have it both

ways. SCO has a clear choice: either pledge not to use any Open Source/Free Software in any of their products, or actively participate in the Open Source/Free Software movement and reap the benefits. For SCO to continue to use Open Source/Free Software while attacking

others for using it is the epitome of hypocrisy."

The latest release candidate is, as ever, available from www.samba.org along with instructions on how to update, and its developers are keen to garner as many user responses as possible to assist further development.

PATENTS AND INTELLECTUAL PROPERTY

How the web was lost?

In a month when many websites fell offline in protest against adoption of software patents, one of its biggest advocates (and holder of thousands of technology-related patents) became the most high profile victim of patent law. The effect on the Internet, commentators have suggested, could be devastating.

Microsoft has been speared to the tune of some US\$520 million in an infringement suit brought by a Chicago-based 'intellectual property' company called EOLAS and the University of California over the concept of applets contained within *MS Internet Explorer*. The patent holders originally claimed a premium of \$3.50 for every copy of IE shipped, which would have led to a \$1.2 billion bill, but the jury decided this would be an excessive charge and dropped it to \$1.50 per user.

While MS mulls its chances of a successful appeal, the likes of Macromedia, Adobe and RealNetworks,

who all develop plugins based on the patent, face the prospect of the most popular browser being stripped of the ability to display much of their content. Meanwhile, projects which use a similar plug in architecture such as *Mozilla*, *Konqueror* and *Opera* – but don't have an MS-sized slush fund – may face similar infringement charges.

The World Wide Web Consortium (W3C) hastily convened a meeting to discuss the problem and prepare responses which, in the first instance, consist of attempting to find an example of 'prior art' (a previous implementation of a similar technology) that would invalidate the claim.

Meanwhile, in the European Parliament, MEPs once again dithered in the face of protests from anti-software patent activists and decided to not decide on their Euro-wide implementation. Despite an attempt to soften the impact with a wording

LINUX PARTNERSHIP

SUSE on SGI big iron



SUSE and SGI have joined forces to put Enterprise Server 8 onto Altix3000 hardware.

Silicon Graphics Inc (SGI) has signed a deal which will see SUSE Linux Enterprise Server 8 installed onto Altix3000-based servers and super-clusters. SUSE, for its part, will offer 'third tier' support to SGI customers and will work with SGI towards a 128 CPU implementation of the distribution next year. SGI claims this will be the first "fully supported 64-processor system running a fully supported, enterprise-grade Linux operating system".

Dave Parry, SGI Senior VP, said the company was attracted by SUSE's support of advanced Linux features, such as scalability, XFS and native 64-

bit environments. "As a highly respected and experienced provider of Open Source solutions, SUSE shares SGI's vision of a broadly scalable and robust Linux environment for technical and scientific computing."

SGI will be site-testing its 128 processor systems with a selection of world-wide partners including the U.S. Naval Research Laboratory in Washington, the Pacific Northwest National Laboratory in Richland, Washington, the University of Queensland in Brisbane, Australia and The Computing Center at Johannes Kepler University in Linz, Austria.



Konqueror and other Open Source developments that use a plug in architecture will be affected by Microsoft's patent suit loss.

change, saying the new rules would cover "patentability of computer-implemented inventions" including mobile phones, household appliances and program-related inventions, the protests haven't dwindled.

The EU was due to make a final decision on September 22.

SCO News

■ SCO has been fined 10,000 Euros in Germany for breaking an order to stop telling everyone that Linux contains stolen code without offering some sort of evidence.

■ The first evidence of the stolen code slipped out of SCO Forum, where Darl McBride presented a series of slides that 'proved' some code had been copied almost word for word from UNIX to Linux. A short while after the leak, the code in question, which dealt with memory allocation, was identified and its original source confirmed. It was previously released under an Open Source license.

■ SCO hinted that it had sent invoices to 1000 prominent Linux users demanding their license fees, but so far, no one has owned up to receiving one. One SCO executive told an Australian newspaper that Linux users would not be targeted – but they would be announcing another 'big hardware vendor' who were infringing their intellectual property (thought to be SGI). He was swiftly contradicted on the question of user targets by McBride.

■ SCO's ownership of Linux is being contested by a paralegal researcher who claims the company has the rights to create an implementation of UNIX. The actual standard – as with the name – resides with The Open Group.

■ Perhaps the most detailed breakdown of all the news coming out of SCO and its rivals is posted on Groklaw – a blog by a seemingly obsessed Linux user. See it at <http://radio.weblogs.com/0120124/>.

Jono Bacon

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



COMMENT

A question of design

“ You know, I like Mac OS X. Yes, I know that some of you may be familiar with some of my (maybe unreasonable?) ranting about Apple computers in the past, but I have turned over a new leaf, given them a go, and I think Mac OS X is pretty funky.

The first thing that grabs people about Mac OS X is its interface. Large colourful icons, subtle animation, smooth fonts and a single mouse button are all associated with the curvaceous beast. Although I admire the graphical quality of the OS, I feel there is a deeper issue we need to delve into.

Apple has done an enviable job of creating an OS that is not only stable and secure, but attractive to both technical and non-technical users. On the new Powerbook that I bought recently, I have running a range of technical and Open Source programs that, in the normal scheme of things, my Father would not even consider using. But when I showed him the Mac, he was impressed with how simple it was to use.

We must remember at this point that BSD forms the base of Mac OS X – as such, we have a means of seeing what is truly possible with an Open Source framework – it can realistically form the basis for an attractive and functional OS.

To me, Mac OS X has shown what is possible with the technology that we all have access to, and although a single point of control – Apple – created Mac OS X, do you think that we can use Linux to come up with something even more impressive for technical users and easy to use for non-technical users?

You know, I think we can. ”



This Christmas, even your technology-fearing in-laws will be asking you "What's all the talk that's been on the TV about this 'Lee-knocks', then?"

PUBLIC AWARENESS

His name is... Linux

Despite its current legal imbroglio, IBM has signalled an increasingly bullish attitude to Linux and Open Source software with the launch of a new 90-second television advertisement designed to promote the strengths of open development and raise the profile of Linux among ordinary computer users, and indeed the general TV-viewing public as a whole.

Set in stark white laboratory-like surroundings, the ad features a small, silent boy with close cropped, almost Nordic, hair being educated by a number of cultural luminaries including film director Penny Marshall, Harvard

Professor Henry Lewis Gates, Sylvia Nasar – author of *A Beautiful Mind*, and legendary boxer Muhammad Ali. Offscreen, two scientists discuss this amazing prodigy as he absorbs such disparate information as the need for the right tools in plumbing and the physics of flight.

The script is designed to highlight the unique features of Open Source, collaborative development. "What he learns we all learn; what he does, we all benefit from."

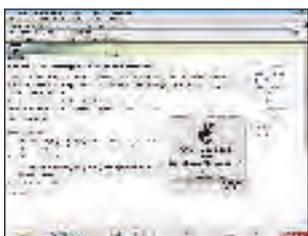
As the ad comes to a close, one of the scientists asks if this remarkable 'Kid' has a name... "His name is Linux."

See [Hoyt Duff's opinion column on the right for comment on this story.](#)

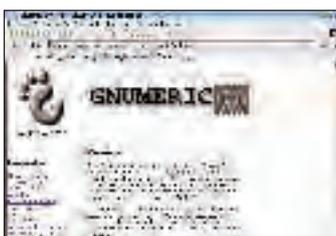


The campaign has a companion website www.ibm.com/open where you can download the advertisement in various formats.

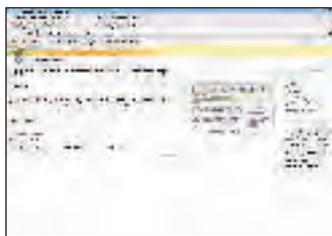
Linux Web Watch/



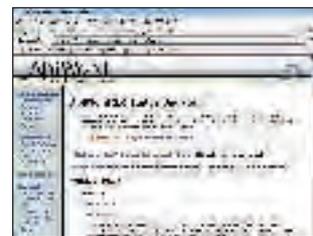
A more refined GNOME desktop.



More Excel than Excel?



GNOME's new browser turns one.



NextGen AbiWord is imminent.

GNOME suite GNOME

A whole clutch of new stable releases of integrated desktop applications

The GNOME desktop (www.gnome.org) project is, at the time of going to press, on the cusp of a new release. As well as the usual graphical and usability tweaks, GNOME 2.4 includes new stable releases of integrated apps.

Gnumeric's most recent update included for the first time compatibility

with all Excel worksheet functions (with a few extra on top) meaning opening .xls documents is far more reliable. (www.gnome.org/projects/gnumeric/index.shtml)

Epiphany, the Mozilla-based application that has supplanted *Galeon* as the default GNOME web browser

achieves its first .0 release. The latest version finds the browser sporting a *Firebird*-style toolbar builder as well as better compliance with GNOME usability guidelines. (<http://epiphany.mozdev.org>)

AbiWord, environment's standard word processor is gearing up for the

release of *Abiword 2.0*. The latest release candidate (1.99.5) includes a number of long-awaited features including table support, headers and footers, revisions and a mail merge facility. There's also a native Mac OS X port for the first time. (<http://www.abisource.com>)

NEWSBYTES

■ Due to 'overwhelming demand' icculus.org has made the first Linux beta of *Medal of Honor: Allied Assault* available for download. To play the game – said to be playable, but with wonky sound – you will first need a licensed copy of the Windows version. Icculus is also doing hosting duties on a number of Linux games including the latest *America's Army* release and *pyDance*, a version of *Dance Dance Revolution* written in Python. *PyDance* is apparently playable using joystick, keyboard or even a dance mat. <http://icculus.org>.

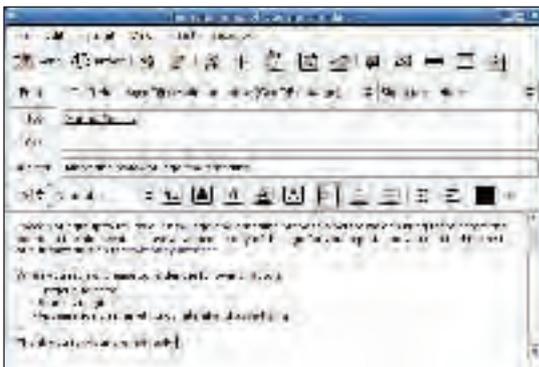
■ US retail giant Walmart has once more expanded its Linux offerings. In addition to a selection of systems pre-loaded with Lycoris Desktop/LX, Lindows, Mandrake and SUSE, Walmart has recently started selling boxed copies of Lindows, Desktop/LX and Win4Lin 5 from Netraverse. A spokesman for Microtel, which builds Walmart's systems, said the decision to stock Win4Lin was especially good for consumers. However, an analyst from IDC says that, though listing on Walmart's site would raise the profile of Linux among consumers, it probably won't drive up sales, which only account for about three per cent of the US market.

■ N2H2 has released its second generation filtering application for Red Hat Linux. The product integrates with many common caching and firewall appliances and provides advanced web filtering tools and simplified administration, as well as a new Virtual Reviewer system which allows users to add sites to the master block and allow lists. www.n2h2.com

■ Microsoft has contacted developers of rival, cross-platform Instant Messaging applications such as *Gaim*, *Trillian* and *Kopete* to inform them of new licensing requirements which will entail either a payment or a lockout. For Open Source projects this will be likely to mean the end of support for MSN.

■ IBM and SUSE have won a contract to bring an archive of more than 80,000 film, TV and video images from the US Library of Congress online. Unusually for a public project, the planners opted for an IBM pSeries solution instead of the tradition cluster of Intel boxes.

■ Following a hoax which claimed Linus Torvalds was running against Arnie for governor of California, self-confessed geek and Open Source advocate Georgy Russell has announced her candidacy. This may not have attracted the popular media's attention had Georgy fitted the computer geek stereotype, but since she's in her mid-twenties and fairly attractive, the famed impartiality of US newsmongers did the rest. www.georgyforgov.com/



Applications such as Ximian's *Evolution* will lead the desktop Linux charge over the next few years.

LINUX MARKET PENETRATION

Desktop numbers game

After extensive testing, Siemens Business Systems claims Linux will quickly climb to the number two position in desktop installations, achieving a penetration of 20 per cent in the corporate market by 2006.

Siemens has been testing Linux on the desktop in enterprise settings with both users and administrators. According to Program Manager Duncan McNutt, the OS is ready for prime time, offering substantial cost efficiencies on large scale deployments of up to 40,000 seats, without sacrificing usability. McNutt said a suite consisting of Ximian's desktop could supplant Windows systems at a similar initial cost, with greater reductions on future expenses, including a 50 per cent drop in hardware costs and savings of up to 30 per cent in administration.

In India, unofficial figures suggest that shipments of Linux systems have surged from almost zero in January to nearly ten per cent in August. Red Hat is gaining most of the benefit from the Indian government's increasing fondness for Open Source software.

Warren Pratt, chief operating officer of SGI, told Pakistan's Daily Times that Linux was 'generating a buzz' especially in India's scientific and defence industries. "Technical, scientific and defence research departments in India are gradually opting to use Linux software," he said. Pratt said two key factors, the country's thriving Linux community and the 'world's second largest pool of scientific manpower' would drive take-up of the OS.

"This is further boosted by the fact that the government of India is encouraging the adoption of Open Source solutions for the country."

However, Microsoft Chairman Bill Gates has said the figures – and those coming out of Thailand – are misleading and that the majority of purchasers of Linux PCs will go on to wipe the disk and install a pirated copy of Windows.



Japan's Trade Minister Takeo Hiranuma announced a tri-nation Linux project.

JOIN THE CLUB

In the wake of the European Union's suggestion that all member states shift towards Open Source software, the governments of Japan, China and South Korea have announced a joint project to develop an international Open Source operating system built along similar lines to Linux.

Initial plans were laid back in March at a meeting of over 100 senior software engineers from government agencies and hardware makers such as Toshiba and Sharp, but the formal announcement of the project was made by Japanese Trade Minister Takeo Hiranuma. The big change from the first announcement, which was concerned with building a rival to Microsoft offerings in the server space, is an added focus on desktop use.

The three partners already have large-scale Linux operations in place, and China's Red Flag Linux has recently signed a co-support deal with HP intended to increase the profile of the distribution at home and abroad. It is rumoured that Red Flag will form the basis of the new OS.

Hoyt Duff

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



COMMENT

IBM: they get it or not?

“ It seems that IBM are in for the long haul and has demonstrated that with the release of an advert themed “The Future is Open”, a nice phrase with ample double meaning. IBM describes Linux as a “...nine year-old boy changing the world”.

It seems that IBM might just “get it” and we would hope that the world follows suit. But wait...

With the usual marketing hype, the ad refers to Linux as “...the most unique innovation operating systems have ever seen.” I don't think so; there are some problems with IBM marketing.

Linux is Unix. Well, perhaps it's more akin to the relationship between Homo Sapiens and Neanderthal, but we'll look past that for the moment. It has a Unix-like heritage that it shares with every commercial Unix, the BSD's and even Mac OS X among others. It's hardly unique, and UNIX has been around for three decades – ancient in OS terms.

What is the unique innovation is the cooperative development model of Open Source. That innovation benefits not just the Linux kernel, but GNU software (and something called the HURD) as well. Indirectly, it has benefited thousands of taxpayers when their governments have adopted Open Source software in lieu of expensive, proprietary licensed software. It's a fundamentally different – dare we say better – way of seeing the world.

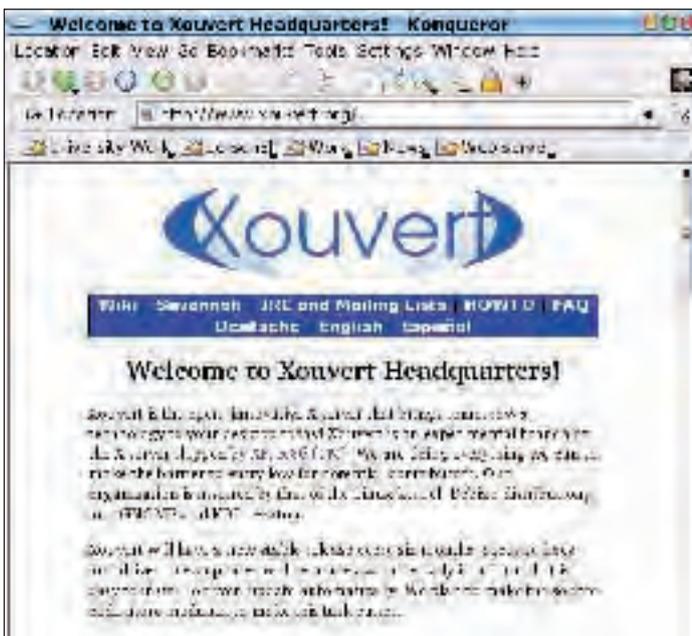
Perhaps even after spending all that money promoting “Linux,” IBM might not ‘get it’ after all. If they do, they might want to share their insight with the seemingly clueless folks in IBM Marketing before it's too late.

AUTOMATED UPDATING?

Replacing Xfree86

Xouvert is a new, forked implementation of Xfree86, that its maintainers claim will put development on a similar footing to successful Open Source projects such as the Linux kernel, GNOME and KDE. Project goals include transparent management, a low entry barrier for contributors and scheduled six-month releases – all of suggesting barely concealed criticism of current Xfree86 management.

The first release, an extraction of the X server source from the standard Xfree86 CVS and compiled with a wide range of video card drivers, is due for release on Oct 1 2003. Subsequent releases will feature big changes including a new modularised source code to make updating and security patching a more fluid – even automated – process. The update facilities will be based on Arch 'changeset management system' created by Tom Lord.



The **Xouvert** (pronounced 'zoo-vaire') project will modularise the X server to improve user maintenance.



The **DLC** will stage a one-day event at Boston University.

LINUX ADVOCACY

Desktop Summit part II

After the last event was 'hi-jacked' by Lindows CEO Michael Roberts, a new group of desktop advocates – the Desktop Linux Consortium – is to stage a rival event on 10th Nov 2003 at Boston University's Corporate Education Center. Supported by HP, Lycoris, SUSE and MandrakeSoft, the event will feature a keynote address by Bruce Perens, who pulled out of the Lindows sponsored show in February.

Perens, executive director of the organisation, said that Linux on the desktop was becoming increasingly viable for many enterprises. "The conference will bring experienced users together with the leading developers of desktop Linux software to provide an invaluable learning environment for anyone considering a deployment of the Linux desktop." www.desktoplinuxconsortium.org

Embedded Linux News

● **Motorola's** first Linux-based mobile phone has received its public debut. The A760 is a traditional clamshell design and runs a selection of Java applications on top of its Linux core. Aimed at the far eastern market, the A760 features a personal information manager (PIM), video and music players, instant messenger and usual phone tools. There's also an integrated digital camera, Bluetooth, Infrared and USB connections. Motorola's decision to embrace Linux is said to be influenced by the Chinese

government's wholesale adoption of the OS. China is not only the source of most of the world's mobile phones, it is also the largest market.

● **Intel** is currently prototyping a 'new concept' in personal networks. The Linux-based Personal Server is a small embedded device containing a hard disk, Bluetooth hardware and a Web-DAV enabled webserver which allows users to read and write data wirelessly from any PC. In other words, HotDesking employees

would be able to carry their profiles, data and personal files from machine to machine in their personal server. The prototype has been built with a 400MHz Xscale processor.

● **G-Mate** (yopy.com) has launched a pair of clamshell Yopy PDAs based on the StrongArm SA1110 CPU. The devices offer a similar feature set, but the YP37000 costs \$499 while the Compact Flash-less YP35000 comes in at \$50 cheaper. www.yopypda.co.uk will have them soon!

NEWSBYTES

■ The overflowing PHP content management system market just got a little more crowded with the release of Rilke CMS. The latest version (0.95) features a number of improvements including a new default theme, single click theme changes, web-based management and a MySQL back end. It's available at www.rilkecms.com.

■ Sony has released a new Linux-based Personal Video Recorder capable of grabbing 342 hours, or two weeks, worth of television on its 500GB hard disk. There will also be a cut down which only boasts a 250GB disk. Unfortunately Sony has no plans to release the devices outside of Japan. Playstation 2 owners can stream media to their console from a Linux, Mac or Windows PC using the \$50 MadCatz Media Player. The system can be used to display MPEG, DivX, MP3, JPEG and other technologies from a wired or wirelessly connected PC to the PS2. www.madcatz.com

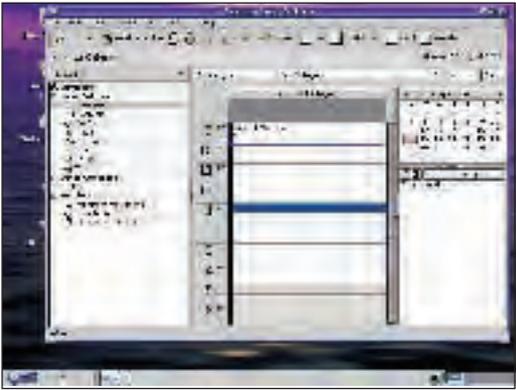
■ The Electronic Frontier Foundation, along with many other organisations, is submitting a petition to US Congress to complain about measures taken by the RIAA against filesharers, including 12-year-old Brianna LaHarra. Concerned? Visit www.eff.org/share/petition/

DESKTOP TRIALS

Migrating in Australia

Telstra, Australia's biggest IT firm, has sent shockwaves through the industry with the announcement of Project Firefly which will see its recently started desktop Linux trial expanded across the entire organisation. Firefly will replace 85 per cent of the company's Windows PCs with a combination of Linux and Citrix-based thin clients. Citing the solid stability of Linux, Telstra's chief information officer Jeff Smith, said the move from Unix and Windows to Linux would slash the group's annual Aus\$1.5 billion IT spend in half within three years.

The initial trial of 250 desktops will be completed by October, and Smith expects the full rollout to take place fairly quickly. Smith also said Telstra would reduce unnecessary duplication by moving systems currently running on HP-UX to Solaris – Sun's support of Linux was the decider – but its vast SAP finance system would eventually migrate from HP-UX to Linux.



Suite will have Exchange-compatible groupware client.



Mad Hatter's browser of choice is Mozilla.

MAD HATTER

Sun toughens up desktop strategy

Sun Microsystems used the chaos caused by the MSBlast Internet worm and the 'end of life-ing' of Windows 98 and NT to highlight the potential benefits of its imminent Linux desktop launch. Citing Sun's own research, marketer Peder Ulander said a recent study of Fortune 2000 CEOs suggested 45 per cent would be looking at alternatives to a Windows upgrade within the next six months. "One of the leading things they're

investigating is Linux on the desktop for certain segments of the business"

As well as highlighting the recurring security problems experienced by Windows users, Ulander said Mad Hatter would 'dramatically' under cut a comparable package from Microsoft. "Whatever Microsoft gives you that special deal at, we'll cut it in half."

Mad Hatter is being built on the same lines as Novell Ximian's Desktop, offering a "fully integrated desktop

stack" including office suite and *SunONE* messaging and groupware applications. It will also demonstrate connectivity to both *Exchange Server* and *Lotus Notes* and, via Wine, will provide facilities to run a number of native Windows applications – an attractive option for many IT managers.

Sun also announced a tie-up with RealNetworks to put the latter's *RealONE* media player onto the Mad Hatter desktop.

PROMOTING LINUX

New online hardware guide

L-Friendly is the latest effort from Michael Roberts' Lindows.com to push the benefits of Linux to the general public – while, of course, publicising the availability of LindowsOS, which was used to test the various bits of kit.

The www.lfriendly.com website launched recently with a small selection of tested hardware broken down into sensible headings such as Linux friendly digital cameras, scanners and ISPs, and a request that users submit their experiences of unlisted hardware.

Roberts is also trying, with some success, to rope 'top tier' hardware makers such as Philips, Sony and Samsung to support the venture.

Epson has a number of products listed on the site, and spokesman Mitch Kadish said their participation was a continuation of their support for Open Source development. "We have been working with the Open Source community as well as Linux distributors to support the printing and scanning needs of Linux users," he

said. "We want consumers to be able to easily find the full range of products that are available and compatible with their systems.

L-Friendly is yet another site, like www.tryoutlinux.com, which at first glance is independent (and much of the information there would be useful to the general user), but on closer inspection has a close affiliation with Roberts and Lindows.



Roberts seems to be creating a huge Lindows-centric network of sites.

LINDOWS PRICE CUTS C-'N'-R

LindowsOS itself has just undergone a price realignment which moves substantially away from the catch-all Click-'N'-Run subscription towards a more pocket-friendly 'buy what you need' system for proprietary applications such as *StarOffice*, while keeping the subscription element for other applications.

The basic monthly Click-'N'-Run subscription now costs \$4.95 or alternatively \$49.95 per year.

David Cartwright

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



COMMENT

Desktop disaster

“ This morning (24 August) I tried to install Macromedia Flash player on my Red Hat 9 box (a fresh, unmolested install except that I'd run the updater to get the latest versions of everything). I downloaded the installer and was a little disappointed that the docs said I had to run it from command line – OK for me, but what about less techie users? I uncompressed the .gz, untarred the tarball changed to the directory with the installer in, and ran the script.

When it asked me for the installation directory, I told it `/usr/lib/mozilla`, because this is what it hinted at, but was told: **WARNING: Please enter a valid installation path.** This seemed weird, because I thought it was valid, so I had a look at the installer script. For some strange reason the script couldn't run unless both `/usr/lib/mozilla/plugins` and `/usr/lib/mozilla/components` existed – and the latter didn't on my setup. I created the directory by hand and the installer ran just fine (and the Flash player seems to work OK). Interestingly it didn't actually put anything in the 'components' directory, it just seemed to want it to be there.

This is dumb, but typical of many software vendors. OK, Macromedia's installer doesn't officially support Red Hat 9, but given that Red Hat is so popular, and version 9 was launched on 31 March – over four months ago – why the hell not? Until vendors like Macromedia get their fingers out, I'm afraid that Linux on the desktop is likely to be limited just to people who want to use (the splendid offerings that are) *Mozilla*, *OpenOffice.org*, (sans *Flash*) and not a lot else. ”

Mailserver

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

★ Letter of the month

This month's winner receives a copy of the O'Reilly book of their choice!



Penguins get airborne!

I was returning from the USA recently, on a new 747-400 aircraft. I noticed the In-Flight Entertainment system had a bunch of new features I had not seen before: Music and video on demand, email and web browsing. I later discovered that the seat-back displays are thin-clients, downloading their software (and media content) from a central server.

However, all was not well. Every seat-back PC was displaying the little penguin, and eventually, as Linux was booting, it repeatedly announced "TFTP timeout..." The TFTP daemon was not running. So, I scribbled a few instructions on a napkin, and gave it to one of the cabin crew. Half-an-hour and one

re-boot later, the seat-back computers were running normally. Later in the flight I had a visit from one of the flight crew; they were extremely grateful – the Virgin Atlantic cabin staff were not relishing the prospect of placating a plane full of passengers, with no entertainment on an 8-hour flight.

Hey Richard Branson, I reckon you owe me a beer!

Simon Rodda, Cornwall

Cool. We were unaware Virgin were using a Linux-based system for entertainment delivery ourselves. It makes a change from the cabin crew asking if there is a doctor/engineer/qualified pilot on the plane! For your outstanding feat of skill at 30,000 feet, we'll be happy to give you this month's star prize. We aim to bring our readers a feature about the Virgin Atlantic's system in a future issue.

Disc directories?

I recently subscribed to your magazine because looking for issues on news stands in North America is pretty much hit-and-miss. And it is usually even more difficult to find a DVD version. I was pleasantly surprised to find out that the subscription service is two months ahead of the issues available on the news stand.

However, please tell me that the omission of the printed listing of the DVD (or CD) contents at the back of the magazine in the September issue was a mistake? The first thing I do when reading a new issue is skip to the DVD content listing at the back to see what goodies are on the new disk.
Tony Hill, Wisconsin, USA

I'm glad you have subscribed and enjoy the benefits that brings you. Sorry you were disappointed not to see the disc listings in the magazine. Over the last couple of months, we have tried to cram so much in that we ran out of space, so something had to give. As the listings are repeated in part on the reverse of the discs themselves, and as there is an HTML front-end to the discs, giving all the pertinent info on the software included, we felt that the space could be better utilised with a few more pages of tutorials.

We don't want to remove parts of the magazine that you feel are useful though. Do any other readers have opinions on the subject? Either email us at the usual address, or discuss it in the 'Magazine' section of our forum

that can be found on our website at www.linuxformat.co.uk

jdiscreport #1!

I am an LXF reader who normally sits on the fence and absorbs the good work of yourself and others. However this time I thought I had to let you know how pleased I was with one of the programs on the July 2003 edition DVD. Obtaining that one simple little program was worth the cost of the mag to me. The program *jdiskreport* is a java version of a Windows program I've been using to manage disk space/usage on my system. The ability to run this program under Linux is extremely useful to me and only makes me more satisfied with Linux.
Neville Cobb, via email

Thanks for the feedback. I was impressed with *jdiskreport* myself. You may be interested to know that an update was released in August, and we'll try to include it on a future *Linux Format* coverdisc.

Choices, choices...

Choice. That's what open source and Linux is supposed to be about. But too much choice can leave users confused and overwhelmed.

I am referring to the large numbers of programs that are out there that seem to duplicate the functionality and behaviour of others. It seems that where two or three programs would suffice for a particular type of software app, other developers then decide that they're not quite right and create

DISTRO WATCH

Small but perfectly formed

Peanut Linux

I installed an old distro of Peanut Linux on a spare machine recently, and was very impressed! I just wondered if you can include the ISO image of the current version – with KDE & MPlayer – on a forthcoming CD or DVD? It's a really good distro and they do not appear to have a UK distributor at the moment. I do not have broadband yet and its a pretty big download even if you have! If this has already been on a DVD then let me know which one.

Bryan, via email

We had a preconfigured version of Peanut Linux 7.5 on the LXF41 DVD. Your request has been passed to our Disc Editor, who will endeavour to include the latest version in a future issue on both the CD and DVD.

For more info on Peanut Linux and downloads of version 9.5, visit www.ibiblio.org/peanut/

Peanut Linux 9.5 supports Kernel 2.4.20, and the whole install only takes 340MB – very handy for running Linux on older machines.



Ogg Vorbis, Shout/ICE Cast, MP3, Wav, Midi, CD, Real Audio, DivX, VOB, all work with Peanut Linux.



That's a lot of functionality crammed into a very small space...

their own versions, all slightly different. For example, to take a particular software genre: web browsers. *Konqueror*, *Netscape*, *Opera*, *Galeon*, *Links*, *Lynx*, *Dillo*, *Atlantis*, *K-Meleon*, and probably many others. Talk about reinventing the wheel! In my opinion, all we need is *Konqueror* (for 'normal' browsing on a 'normal' PC), *Galeon* (for lower-spec machines) and maybe *Lynx* (for text-mode browsing). *Opera* may get a look-in as well.

After we put all the other superfluous programs to one side, we'll have freed up many developers who can work on making the remaining programs in each software genre really good indeed – and the more powerful the apps that we have on Linux, the greater the incentive to jump platforms. DTP programs for example are pretty dire on Linux – so let's club together and change that! Basically to all developers out there: are your programs actually

necessary, will they actually be used, and could your coding efforts be better spent on other projects?

M Wise, Croydon, Surrey

This has been the age-old debate about Linux, and indeed the Open Source movement since time immemorial. Some people relish the fact that they can choose – choice is what freedom is all about, after all. Others can certainly

find themselves confused by a multitude of possibilities.

The problem with choosing one browser, for example, is that not everyone is going to agree on what that one should be. I dare say the *Galeon* development team aren't going to throw up their hands and say "Well, *Konqueror* is much better, let's go and do something else instead", for example.



Excellent documentation on all aspects of Red Hat Package Management is available from the <http://freshrpms.net> homepage.

Posting to the forum

The LXF online community

Got a burning Linux issue you want to discuss with other LXF readers? 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 covering many different subjects. Give a hoot – contribute!

Some Linux distributions have taken this viewpoint though. Take a look at Lycoris – the menus aren't filled with the possibly confusing names of various different apps, but are task oriented. The distribution is deliberately slimmed down to prevent new users getting confused by too many applications that perform similar functions. For the purposes of developing good software, new ideas to circumvent old problems are often generated by those projects which you might not consider to be the best examples.

RPMS

Many thanks for the Package Management tutorial in LXF44. I've only been using Linux since January this year, and have yet to experience much "dependency hell", but I do find the supplied GUI tool more than a little lacking in granular package selection.

I was very concerned with the caption under Fig2, however where you state "kernel updates... should never be installed via up2date". I would agree it's probably not a good idea to let up2date upgrade the kernel automatically especially if there are likely to be driver issues.

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

MAILSERVER

« I've updated this way at least five times since installing and never had a problem other than having to reinstall nVIDIA and Atmel driver modules. You do realise that the process leaves your existing kernel intact, including the *grub* menu option? And that you can simply remove the old kernel (or the new one) when you wish with the appropriate `rpm -e` command?

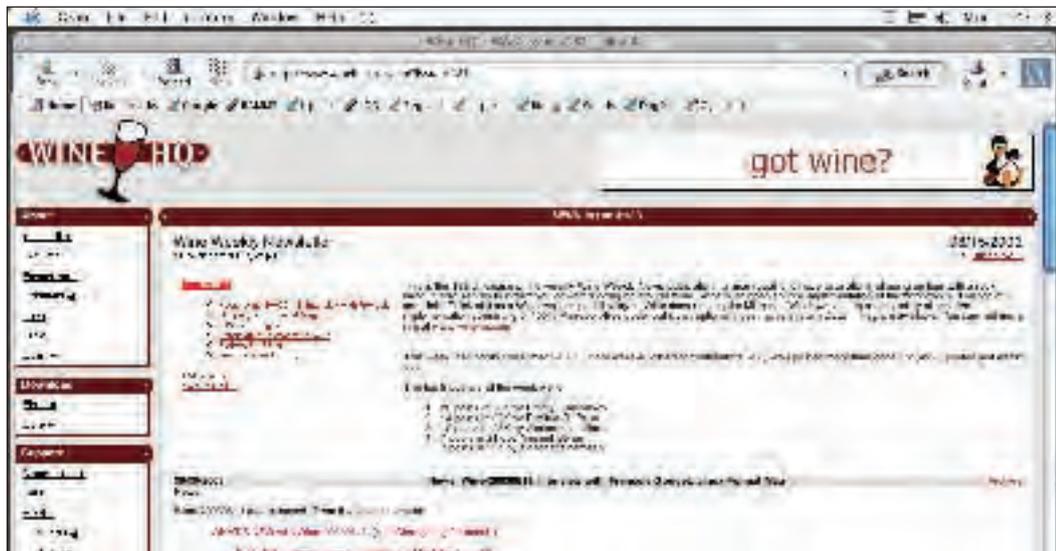
Is there some other factor I've missed that puts my system at risk?

David W. Haggett, *via email*

Really, the advice given there depends of what you use your system for, and falls into the category of "if it ain't broke..."

Wine

I happened across the latest Wine newsletter and I thought I'd drop you a line in case you missed it – your fame spreadeth! www.winehq.com/site?issue=183 Keep up the good work guys, as I think the magazine is



Linux Format's opinion is valued throughout the Linux community – see *LXF42* for *All About Wine*.

excellent and a major change from all the clone mags. It has the right balance of basic and advanced techniques for both newbies and old

hands alike. As I am a Windows administrator, I need to see both 'sides', though sometimes the basic articles can be too basic, yet even though the concept is 'obvious' there is generally something you point out that I had missed. I have also found the current articles on *Samba* invaluable as I have just deployed a test machine in my company. I needed *Samba* to facilitate authentication to a *Squid* proxy. (Any chance of a guide to using *Squid*, or have you done that one recently?)
Neil Williams, *via email*

We love compliments – keep 'em coming! We reviewed *Squid 2.5* back in *LXF37* – order your backissues on page 101. By way of an explanation, *Squid* is a very popular proxy server that runs on Linux, FreeBSD, Mac OS X, most flavours of Unix and also Windows NT, so it's obviously a very important part of the Open Source world. When 3.0 finishes its development cycle and becomes stable, we'll probably review it, and if there's enough new functionality included to warrant another tutorial, *LXF* will provide one of those too. [LXF](#)

“In developing new software, new ideas to circumvent old problems are often generated by projects which you don't consider to be the best examples...”

Helpdex

shane_collinge@yahoo.com



READER TIPS

MODEM COMPATIBILITY

It was frustrating to see the reply to Gerry Green in *LXF43 Answers*. He wanted to be advised of a Linux-compatible internal modem.

I was asking the same question a couple of years ago and usually got the same "answer" – a lecture on Winmodems, and "you need a modem with a controller [insert arm waving here], but get an external one, it's easier". But, infuriatingly, no one ever seems to give actual examples of internal modems with such 'controllers'.

But I was building a PC from scratch and wanted no compromise, and no clutter of subsidiary boxes and interconnectors. PCs have expansion slots for this very reason. Hard drives and network cards were taken inside the case years ago, so why are modems so slow to follow? Yes, most internal modems are Winmodems unless they are old ISA ones, but 56K PCI Linux compatible ones exist and it is no sweat to list them, being so few. The compatible ones that I know of are:

Multitech MT5634ZPX-PCI
Zoom 2920
ActionTec PM560LKi
IBM 33L4618

These are up-market as workstation modems go, and reassuringly expensive – say three times the cost of a Winmodem. I use the Multitech – it has no hardware lights to watch (software ones though), but it just gets on with it.

Robert Phipps, *via email*

MANDRAKE INSTALL

On page 9, *LXF44*, Richard complained that the Mandrake 9.1 installation, graphical interface configuration, "hangs at the test screen". I experienced the same problem (or at least the same symptoms) on a similar setup. After several failed attempts, I chose 640x480 and a lower colour depth and did NOT test it.

The installation process continued fine and the system rebooted. I then simply changed things back to the desired settings in the control center.

A bit hit-it-and-hope I know, but it worked fine for me!
Matthew, *via email*

SD_MOD LOADING

I have just got around to reading *LXF42* and notice that Ralph Cox was asking about missing /dev/sda entires under Mandrake 9.0.

I had the same problem and discovered that all that is required is the `sd_mod` module which is not loaded automatically even if usb-storage gets loaded. Adding a line with just `sd_mod` on it to /etc/modules will ensure it is loaded at boot time. Hope that helps Ralph and your readers!
Paul Fotheringham, *via email*

USE FOR WINDOWS

The reason I wrote was to add to the letter *Scanner episode II* on page 100 in *LXF44*. I have had the same problem with my Mustek 1200 plus scanner which shows up as a 'Bearpaw'; and an HP

scanner which shows up the same. I have compiled the software (*sane*) with the correct firmware and used *rpm* on Mandrake 9.1, 9.0 or 8.2 and the same thing happened I kept getting the error `gt68xx:/dev/usb/scanner0:invalid argument`

But I can get it to work, all I have to do is boot into windows first then restart (not shutdown) and straight into Linux and the scanner works fine. (So there you are, a use for Windows!) I have no idea how this works, but it does. I think this works with Slackware as well. If anyone knows *why* this works, please share it on these pages!
Mitch, *via email*

NEWPKI 2.0

I have nothing to do with it myself but I found this an impressive Open Source project. It's sure a lot more interesting than yet another window manager. For more info, see www.newpki.org. The license of the version 2.0 of *NewPKI* changed, it's now fully GPL.

Some of the main changes from earlier versions are:

- Handling multiples PKIs per physical architecture
- Handling multiples CA chains per PKI
- Central configuration of the whole PKI
- Handling the offline CA mode
- Smart Repositories, to handle the PKI Objects (Entities' configuration, CSR, CERT, CRL, ...)
- The PKI can be deployed on multiple sites, with automatic synchronisations
- Handling custom 'plugins' for the Publication and the requests validation on CA side
- Improved mail notification, with the possibility to activate audits
- Handling groups of users
- Handling ACL per users and per group (delegation of roles to users and groups)
- Improved logs display with filters
- A very simple way to put entities in relation with each others
- Improved security (all sensitive datas are crypted and signed)

The release version of *NewPKI v2* is expected in September.

Although this version is

incomplete, it should give you a very good taste of what the final version will be, here the list of the present entities:

- The *NewPKI* Server
- The PKI Entity
- The CA Entity
- The Repository Entity
- That's enough to be able to create CA Chains, and sign PKCS#10 CSRs directly from the CA GUI. The missing entities have just been published along with new betas over the summer, here is a list of those missing entities:
- The RA Entity
- The Publication Entity
- The Key Escrow (Store) Entity
- The End Entity

Hope your readers find this of use!

Bart Symons, *via email*



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.



5-4 Average.



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

THE TOP STUFF AWARD

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



WHAT'S NEW...

Armari Pro3D workstation >>

64-bit Opteron chips aren't just revolutionising the server market – see just why this neat box merits a high score **p18**

Fujitsu Primergy TX300 server

If you haven't got a big rack to put your servers in, this tower option wouldn't look out of place in the office **p20**

OpenOffice.org 1.1

1.0 just added stability rather than new functionality, but 1.1's Changelog reads like a wishlist for feature-junkies – now's the time to upgrade! **p22**

Aethera

OpenGroupware disappointed us last issue – will TheKompany's attempt at an *MS Exchange*-alike be any better? **p24**



Linux peripherals

NEW SERIES! Many readers have requested this – over the next few issues, we'll be examining hardware that is 100 per cent compatible with most Linux distros **p26**

Book reviews

Algorithms for Compiler Design; Linux in a Nutshell; Managing Linux Systems with Webmin; Firewalls for Dummies; Operating Systems Fundamentals **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 and PHP 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**



OPTERON-BASED WORKSTATION

Armari Pro3D Workstation

Opterons aren't just server chips – 64-bit processing is handy wherever number crunching is key, like the workstations of today. **Paul Hudson** takes Armari's new system out for a test drive...

BUYER INFO

64-bit workstation – the first of its type that *LXF* has reviewed – expect the competition to run for cover...

- **DEVELOPER** Armari
- **PRICE** £4,000 (inc. VAT)
- **WEB** www.armari.co.uk

Some have said in the past, “God gave us servers so we could run Linux”, emphasising the fact that Linux has always run best on high-end business machines that rarely had to interact with users. The Linux desktop movement has been going for some time, trying to get Linux more widely accepted by users looking to replace their copy of Windows 98 with something else, and, thanks mainly to the quality of products such as KDE and *OpenOffice.org*, Linux now has a very firm foothold on home machines – not as strong as its foothold in the server market, but a good start nonetheless.

One arena that has yet to see little activity on the Linux front, however, is the workstation market – high-spec machines designed for the desktop, but with the power of servers, designed to perform complex business tasks such as CAD/CAM, graphic and video processing, and local database handling. However, that's all set to change thanks to Armari...

Little black box

The Armari machine, from the outside at least, could be from anywhere. Granted, the black colour and Armari logo stickers signal that there's more than meets the eye, but few people

would correctly guess quite *how much* more there is inside – this is a machine that should, on the basis of numbers alone, keep up with or perhaps even exceed the servers we usually test in the magazine.

At the core of the box are two of AMD's latest 246 Opteron CPUs, running at 2.0GHz apiece. These have only been launched very recently – which itself underlines how close Armari is to AMD. The CPUs are backed up with 2GB of ECC RAM, a 120GB Western Digital UDMA/100 hard drive, and a Plextor PlexWriter 52x CD-RW drive. Sound good? You've not heard the half of it – the graphics card powering this beast is a 128MB nVIDIA QuadroFX 500 card running on a 8x AGP interface. This is a very well-renowned enterprise-level card, certified for use with a variety of graphics applications, and theoretically able to transfer just shy of 8GB/sec through the graphics memory.

With a graphics card like that, it's no wonder the monitor supplied is also first-class – the machine comes bundled with a 19-inch Samsung TFT LCD monitor, capable of happily supplying a resolution of 1280x1024 at a 500:1 contrast ratio. Graphically speaking, this is pretty much as good



as it gets right now, although users it's possible the most demanding users might want to try one of the more expensive QuadroFX cards.

This is clearly a dream machine for just about everyone out there, and it's clear that Armari hasn't spared any luxuries across the board. Even the supplied mouse and keyboard are Logitech's new cordless desktop MX combo, the model up from the kit reviewed on page 26, and features the new rechargeable “mouse house” dock.

Performance, though, is of course what counts to most people – does this machine justify its £4,000 price tag on the basis of performance alone?

Go faster stripes

On paper, this machine should quite happily outpace everything else we've

reviewed so far, thanks largely to the processors driving it. Granted it only has a single hard drive, and that's connected using IDE as opposed to SCSI, however it's a particularly high-quality drive, and should still perform well. It's important to remember that hard drive access transfer rate is much less important on workstations than on servers, simply because the tasks that are performed are often less disk-based.

The star performer is undoubtedly the *oggenic* score of 6.18, which is much faster than we thought and shows that the performance of this machine isn't simply down to it having the latest CPU. Armari are experts in the field of computationally intensive tasks, and they have obviously put their experience to very good use in the machine – you can expect all CPU-intensive tasks, from video processing to audio encoding to perform marvellously.



Apache benchmark is also fairly predictable, as machines this fast simply can't help but serve web pages at a lightning pace.

Once you take into the account that SLES does seem to have a serious problem with either MySQL or PHP, causing that benchmark to really grind even on this hardware, it's apparent that there's very little that can hold this hardware back from greatness. Yes, £4,000 is a lot of cash, but not when you consider what you get for it, and especially not when you remember that it all comes bundled with a first-class graphics card and monitor. Furthermore, considering that companies are the primary market for this product, the price works out at only £3,399 once VAT is discounted, which makes this machine an absolute bargain.

Noise and equipment

One of the most commonly cited downsides to Opteron's is that they tend to run quite hot, therefore requiring some fairly extreme fan cooling. This machine, however, features SuperMicro's new low-noise chassis, and also features a dual-fan, low-noise PSU, with the end result that the machine is no louder than an

average desktop box – quite an achievement for Armari. With a noise level as low as this, and such an attractive mouse, keyboard, and monitor bundle, this machine would fit in well on any corporate desktop, or even on a desk at home.

One minor downside is that there was very little else in the box – we would at least have liked to have seen SLES CDs, or a warranty sheet. Furthermore, it's an odd choice not to include a DVD drive given that this machine is able to do so much on other fronts. The lack of a DVD ROM isn't too bad, because at least Armari lets you configure the parts yourself before the machine is assembled. On that note, one advantage to this system is that you can configure *all* the parts of the machine before it's shipped – processor speed, RAM, hard drive *etc.*, meaning that you can change the system over to SCSI if you want maximum hard disk performance, and you can even upgrade the QuadroFX to the best in its class to squeeze every last drop of graphical goodness out of the system.

Bang for buck

It's clear that this is a machine designed to make the most cutting-edge hardware available outside of the server room, and it does that remarkably well. The new Opteron 246 CPUs have scaled perfectly, showing that we're only at the very beginning of AMD's 64-bit adventure, and also that Armari clearly work very closely with AMD's engineers to make sure that all the hardware works in perfect unison in order to get the most out of the chips.

Armari's history of designing and producing servers dedicated to high-performance computing (HPC) has clearly biased their workstation design more than just a little: yes it comes in a nice little chassis with a great monitor, but internally there's a beast just waiting to get out if you give it some work it can sink its teeth into.

For use as a workstation, a lot of effort has gone into make the machine extra quiet, and the results are quite impressive – even during the hard drive benchmark, when the disk was being thrashed by our tests, the machine was still virtually silent. The desktop machine being used to type this review is louder than this Armari workstation, which should really say it all! We think the price of this machine

is wholly justified by the exceptionally high build quality of the machine, as well as the high calibre of peripherals bundled with it. That said, a DVD ROM would top it all off perfectly.

The fly in the ointment, sadly, is SUSE Linux Enterprise Server. For some reason, SLES has serious problems running the MySQL/PHP benchmark that don't exist in other distros, and we're at a loss here at LXF as to what we can do to work around the problem. When tested with Mandrake, the system was almost ten times quicker, despite it not being able to take advantage of the Opteron's 64-bit power. Granted, few workstations will want to thrash the system quite like we do, but the test shows that there's a problem somewhere in the system and that may well have adverse affects on other parts of the system also – sort it out, SUSE!

Despite the glitch with the OS, the Armari workstation still manages to fight through and earn a respectable score of 3.18 in our benchmarks, which is very good for a machine of this size and price. The only item holding the system back from greatness is the IDE hard drive, which was the same problem we found with the Armari server from last issue, when we said "we'd be very interested in seeing a machine that has a better hard drive (preferably with RAID), as the chances are that's the only thing holding this thing back from a score like we've never seen before". QED! **LXF**

Sadly, the high *oggenc* score is dragged back by another poor MySQL score, scoring just 0.3 when running on SLES. This is a problem we have encountered before, and the blame unequivocally lies with SUSE – the RackSaver machines we reviewed back in LXF43 had the same problem, but then scored a respectable 2.99 when running Red Hat. Similarly, we put Mandrake 9.1 on this machine and tested it again, and it scored 2.5. It's rather annoying that SUSE has yet to fix this problem, or even acknowledge that a fix is needed, but it's important you understand this is *not* a hardware issue because the hardware works excellently under other distros.

The kernel compilation score is pretty much exactly as you'd expect looking at the scores of other Opteron's we've reviewed – here the score continues to scale linearly. The

VERDICT

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

Desktop supernova with heavenly prices and stellar performance; Armari grabs the lead in the Linux workstation market with a world-beating debut.

LINUX FORMAT RATING
10/10

BENCHMARKS

hd	1.45
apache	3.5
mysql	2.5
compile	2.25
oggenc	6.18
Overall	3.18

SERVER

Fujitsu Primergy TX300

It's not often we get tower servers in for review, mostly because they won't fit under Paul Hudson's desk. Nevertheless he's got the latest from Fujitsu – but does it beat the competition?

BUYER INFO

Powerful general workplace server. For extreme (but noisy) performance see servers on graph (opposite page)

- **DEVELOPER** Fujitsu
- **PRICE** £3990 (inc VAT)
- **WEB** www.fujitsu-siemens.co.uk

Despite Fujitsu Siemens being one of the largest computer manufacturers in the world, we hear surprisingly little from it. That's not because it doesn't like Linux, far from it, but simply because it's been keeping its offerings very low-key. When this Primergy box arrived on our desks, amid much heaving and grunting thanks to the sheer weight of the unit, we all had just one question on our minds: is this the best-kept Linux secret?

Here we are now

Let us lay the facts bare: it's a dual 3.06GHz Xeon tower server with 3GB RAM, 4x36GB hard drives set up as a stripe set, and hyperthreading enabled. It came supplied with Red Hat 8 preinstalled and configured.

The unit itself, as it isn't a rack server, has actually been designed to be aesthetically pleasing – it has a sliding metal grill front to cover the RAID system, but also a DVD drive, a floppy drive, and, wait for it, a CD-RW drive too. This is the first server we've ever reviewed that comes with a DVD drive, and it's also the first that had a CD-RW also – an interesting combination of hardware, but it's certainly most welcome.

Around the back there are two USB slots, two gigabit NICs, a redundant PSU, as well as the usual ports. There is an extra USB port on the front, so this is definitely a modern server well able to handle modern tasks.

There are three particular points of note about this machine, and those are its size, weight, and noise level.

Firstly, being a tower server, it's shaped something along the lines of tower desktops, with the difference that it's about twice as long. It's hard to say what takes up all the space inside, but at the very least it provides a lot of air-cooling for the machine. The weight of the box is also interesting, as it really feels like Fujitsu have filled the bottom of the machine with concrete – it's incredibly heavy for such a small unit.

On the noise front, this actually makes quite a change from pretty much every machine we have reviewed recently. If you've been keeping close watch on our server reviews over the past few issues, you'll have noticed that the boxes' general noise level seems to have been reaching a crescendo – as more and more power is put into servers, they seem to need more and more fans. Not so with this machine: immediately after powering on, yes, the fans came on strong and loud. However, the noise level dropped drastically, and is about equal to one loud desktop PC, or about two average ones. Naturally this makes sense given that this is a 'normal' server, and thus is subject to much tighter restrictions than a rack machine would be.

Entertain us

With 6GHz of computing power behind a single server, as well as the fact that it has an excellent hard drive configuration, it's hard not to have high hopes for its performance. With a lot of RAM and Intel's leading CPUs powering this thing, it's clear that it should perform best at CPU-related tasks – this is reflected neatly in its oggenc score of 4.68, which is amongst the highest we've seen. This is backed up with a very strong MySQL result of 3.5, which is the best we've seen by quite a margin.

It's important to remember that this isn't a rack machine – it's designed to work as a departmental server, and therefore has a lot of CPU capacity;



work like serving MySQL queries is its bread and butter. With this kind of excellent performance, it's unsurprising that the Primergy recently hammered HP to take the top slot in TPC results for the two-way server segment. The fast system speed carries over into the kernel compilation results – 2.16 is another record score for us, which makes this a very mean performance system across the board.

The other scores are quite good, but don't pull away from the pack by

any large margin. We were a little disappointed by the hard drive score, particularly given that this machine has four drives chained together. Yes, this will see a drastic improvement as soon as U320 SCSI is implemented, which will improve a few of the other scores also. As long as quality hardware is used, which is almost a given considering this is Fujitsu Siemens, we expect the hard drive score to be closer to 4 than 3. Once you add in the small speed boosts this will provide in other areas, we expect

this machine will score around 3.4 in the benchmarks overall, which is quite a healthy improvement at no cost.

It's fun to lose

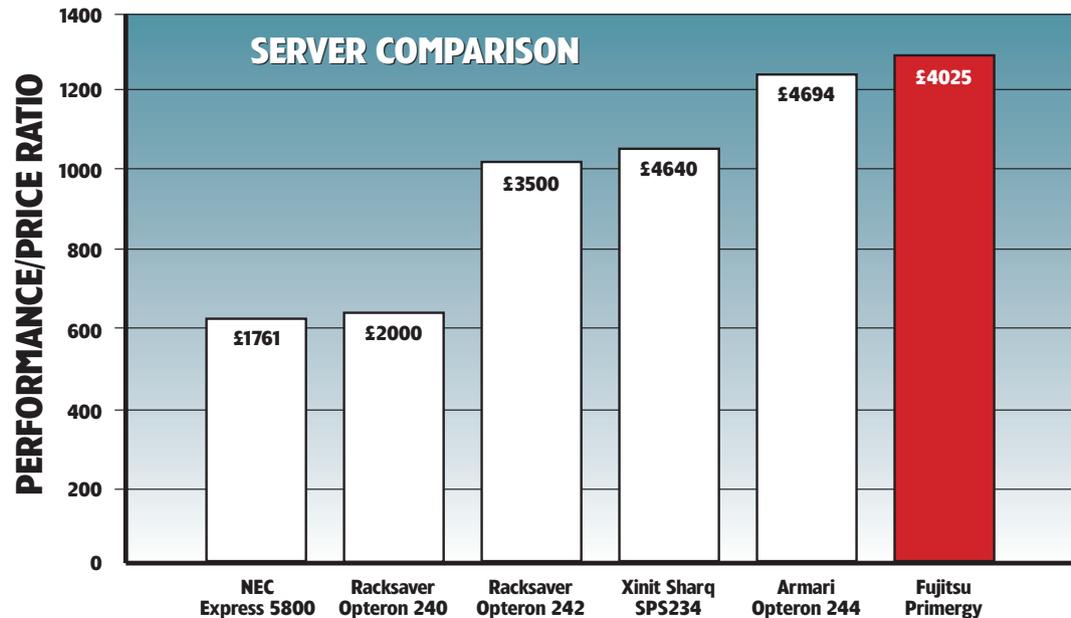
While the machine sees several records broken on our side, and runs very quietly for such a monster piece of hardware, it's a long way from perfect. The biggest gripe was that the machine, a dual 3.06GHz Xeon with 3GB RAM and four hard disks, was shipped with Red Hat 8. Now, while you and I might groan when we see that various companies, such as Red Hat and SuSE, sell features the stock kernel already has and brand them as "Enterprise" despite normal distros like Red Hat 8 having them too, that's not the problem here – we realise that Red Hat 8 is a pretty good distro with lots of power behind it, and it can easily handle this hardware. The problem is the *support* – both Red Hat and SUSE sell their software based upon the quality of support backing it up, and the level of support behind Red Hat 8 scores a big, fat zero on the LXF Supportometer.

To make matters worse, the RH8 install that was shipped was half complete – we didn't have GCC or various other key components, and, annoyingly, we weren't given any CDs either. Unsurprisingly we took the default Red Hat install off, and tried various other distros out – Debian, Mandrake, and SUSE all work fine on the hardware, which is good to see.

One particular high point is the performance/price ratio of this machine. If you look at the graph above you'll see that when you divide performance by price, the Primergy (coloured red) gets a score of 1294. Bearing in mind that higher is better, the Fujitsu actually comes first amongst the other machines – a very impressive place indeed.

Another brick in the wall?

We review a lot of servers in this magazine, mostly because that's where Linux holds the majority of its market share. The difference



between most of the reviews we do is usually speed (and noise!), because as we get sent newer and better machines the scores go up and up.

With the Primergy, however, things are quite different. Yes, there's absolutely no doubt that it's quite a fast machine, but it's also very easy on the eye, and also makes less noise than a desk fan. To add to that, by the time you read this review, U320 SCSI controllers and drives will already be sold as standard in the machines. Add to that the fact that this machine comes top when you consider performance/price, and it seems that this really *is* the best-kept Linux secret.

On the downside, this machine doesn't have the highest performance if space is your chief concern – if you're in the market for raw performance in a tiny space, this is the wrong place to be looking, because this machine is designed for more than just blade use. Aesthetically speaking, this is a very attractive server – we love the fact that it's quiet,

well-covered for ports, comes with practical conveniences that modern admins expect – such as a DVD drive, and still comes wrapped up in an attractive case that's actually had some thought put into its design. Moreover the server came with lots of bundled CDs containing a great selection of documentation, support tools, and more.

If you're buying a server purely for raw computational power, you'll almost certainly be better off with an Opteron. However, if you're looking to change your IT department over to Linux, perhaps even using SUSE's new Standard Server, this is the best choice we've seen so far. [LXF](#)

BENCHMARKS

hd	2.81
apache	2.4
mysql	3.5
compile	2.16
oggenc	4.68
Overall	3.11

VERDICT

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

Top-notch configuration, clever design, lightning performance, and you could easily live with sitting next to it all day.

LINUX FORMAT RATING

9/10

PRODUCTIVITY SUITE

OpenOffice.org 1.1



The last major release of OpenOffice.org, v1.0, was almost eighteen months ago – is this latest release going to continue to drive Linux's most popular office suite? Paul Hudson seems to think so...

BUYER INFO

The definitive Linux office suite just got even better, including additional MS Office compatibility!

- **DEVELOPER** OpenOffice.org
- **PRICE** Free
- **WEB** www.openoffice.org

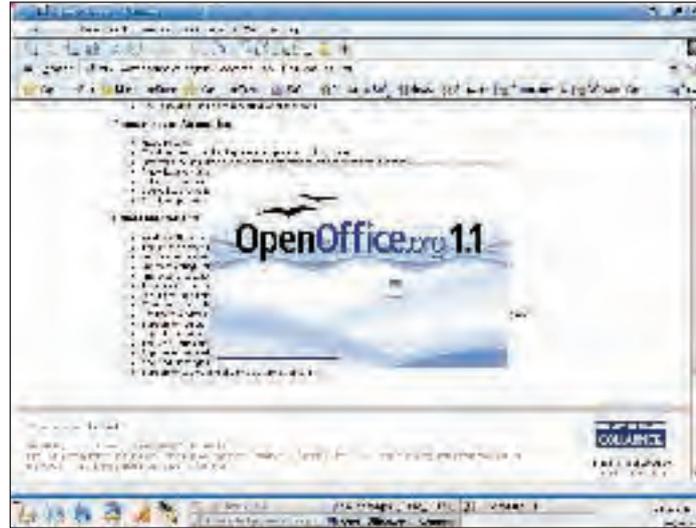
With the first release of OOo, the focus was getting a stable product out the door, and to a large extent they were successful – OOo is available on every desktop Linux distro worth its salt. With 1.1, the development team has taken the opportunity to focus new features, and after an extensive beta-testing cycle the final version is now available.

Import and export

The various point releases of OOo1.0 mainly addressed bugs in the product, with new features being almost non-existent. While this didn't garner support amongst feature-junkies, most will agree it was important for OOo to get the reputation as being a mature and stable product. The 1.1 release was the first time the developers had the chance to break away from that mold and try out new features, and they really went to town – the Changelog reads like a wish list, and users upgrading from 1.0 will really get the 'new car' feeling.

Perhaps the most popular new feature will be the PDF export filter, which makes generation a PDF file of your documents as easy as clicking 'Save as PDF' – this will be a great selling point to consumers who have had to create PDFs regularly using proprietary software until now. Also on the import/export front, it's now just as easy to export *Presenter* documents direct to *Macromedia Flash*, which is another feature you just won't find in competing packages.

Those of us who have to interact with Windows and Mac colleagues who use *Microsoft Office* a great deal



It's surprising how much difference a progress bar makes to the perception of speed by users!

will be pleased to hear that the already excellent import/export Office filters have been improved even further – *Writer* now supports forms conversion, *Calc* now supports data validation from *Excel*, and *Presenter* has noticeably more accurate *WordArt* object translation.

Macros and stability

While importing and exporting are what will wow the masses, the features we found most useful were just sitting in the background. For example, OOo 1.1 now comes with a fully working macro recorder, that allows you to record and playback actions executed within the program – a feature present in tools like *MS Word* for some time, and very welcome in OOo.

The latest rash of big new features have also introduced a variety of big new bugs – most (all?) of which have now been fixed, of course, but they also inspired two interesting new features: crash reporting and faulty file loading. We've all seen crash reporting in operation in other tools, and it works much the same in OOo – when it bails out, it can collect a variety of information about your machine and the crash to help fix problems. Faulty

file loading is an excellent feature that will automatically try to recover slightly damaged OOo files – time will tell how well this will work.

Usability and speed

The help files, notorious for being hard to navigate around, have received a welcome face lift and update, which, along with various other usability improvements such as detecting misconfigured Java setups, automatically cascading new windows, and autodetection of newly installed spelling and hyphenation languages (this caused much pain in OOo1.0), make the new release much easier to use for people who just want to use an office suite without hassle.

While all these new features are very welcome, it's good to see that a lot of work has gone into speeding up the suite – both materially and immaterially. That is, the whole suite is a great deal faster in loading programs and also loading documents, but a variety of small tweaks have been implemented to make things feel faster even if they aren't. For example, the splash screen now includes a progress bar so you at least know that it's doing something – a tiny difference, to be

sure, but a noticeable one. We tested OOo on an 800MHz PIII: *Writer* starts in just under 8 seconds from scratch, and about 1.5 seconds if an existing *Writer* is already running.

The OOo standard

OpenOffice.org is now the standard office suite on Linux, which gives it an enormous amount of power, and also responsibility, in the open source world. If you're a *KOffice* user currently thinking "*KOffice* is much better!", don't get me wrong – *KOffice* is a great office suite, but there's a *good reason* why the *KOffice* developers intend on switching all *KOffice 1.4* applications over to using the *OpenOffice.org* file types.

Many users may be tempted to look at OOo1.1 and think "I'll wait for 1.1.1, because there are probably many bugs", but the reality is quite different – if anything, 1.1 is a great deal *more* stable than 1.0.3.1, the previous stable build, and is highly recommended in terms of reliability if nothing else.

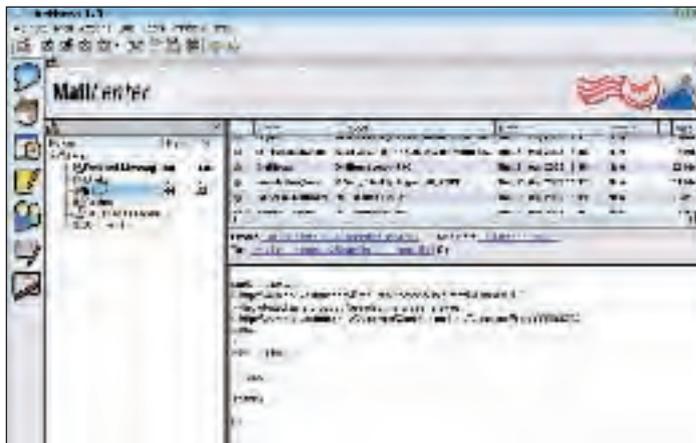
When considering the excellent new features, improved speed, added reliability, superior import and export, and enhanced documentation, you'd be forgiven for wondering why this release isn't called version 2.0. OOo1.0 was a great product, but even that can't hold a candle up to this great successor, so congratulations are in order for the developers, and to all you users – what are you waiting for? Upgrade now! **LXF**

VERDICT

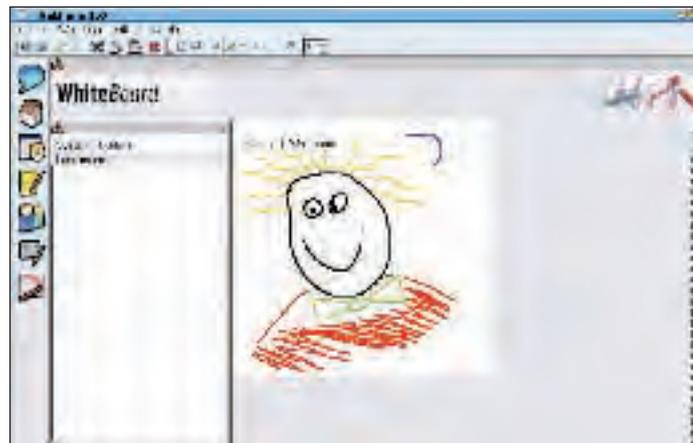
Features	10/10
Performance	9/10
Ease of use	9/10
Value for money	9/10

Faster, better, stabler, and lots more adjectives that mean 'pure class' – OOo1.1 is sure to be a hit.

LINUX FORMAT RATING
9/10



Mail client had a few random crashes, but is better than *OpenGroupware.org*.



The *Whiteboard* plugin, again, is simple but effective.

COLLABORATIVE WORKING

Aethera 1.0

Building on *KOrganizer*, TheKCompany has launched a new groupware/communication suite. **Andy Channelle** reorganises his life.

BUYER INFO

Mail, calendar and groupware solution. Alternatives include *Ximian Evolution* and KDE's nascent *Kontact* project.

- **SUPPLIER** TheKCompany
- **PRICE** Free (Whiteboard and IM plugins: \$9.95)
- **WEB** www.thekompany.com

It sounds like one of Godzilla's arch-enemies or a heart condition, but is in fact a 'new' messaging suite from TheKCompany. Built around a core of KDE's personal information manager application, *Aethera* adds an *Evolution/Outlook* style user interface and plugin architecture to create a single app which should, in theory, be able to handle your day-to-day communication – email, contacts, to-do lists, schedules etc.

The key competitors here are the fairly mature *Evolution* from Ximian/Novell and the newly minted solution from KDE's *Kroupware* effort, both of which offer a similar featureset. But, to stand out, TheKCompany has a pair of proprietary plugins at \$10 a throw which extend the functionality of the suite to include *Jabber* instant messaging and a whiteboard.

The *Jabber* plugin allows users to chat with those running MSN (for now), Yahoo, ICQ, AIM or another *Jabber* client and a future version will

provide file transfer between users. Meanwhile, the Whiteboard plugin also uses *Jabber's* messaging protocol for online collaboration.

Installation was easy on an RPM system, though it does need a small patch for font anti-aliasing, but a colleague running Debian suffered serious library dependency problems with both a converted RPM and tarball. Once the app is installed, the two plugins from TheKCompany were equally simple to install; though a menu entry would have been nice.

On firing up for the first time, you're presented with a fairly uncluttered interface with icons representing installed plug ins ranged down the left. Each bit of the interface can be pulled out to become a floating menu, but I found reattaching them hit and miss and still really have no idea how I got the application back to its virgin state!

Day-to-day

With the right details to hand – names of incoming and outgoing server etc – setting up *Aethera's* mail server is a breeze. The mail client itself is the usual three-paned affair with folders on the left and the main window split between mail headers and a preview. My first attempt at connecting to my mail server didn't bode well though. Connection and authentication went fine, but mid-download the application

disappeared. Oh well, try again; same result. I ran through the mail boxes one by one and narrowed the search down to a Blueyonder account. One unremarkable mail in that account made the whole application fail, and since deleting it, *Aethera* has been admirably stable. There are a few missing 'luxuries' that users of *Kmail* and other clients may have grown used to: audio notification of incoming mail, plain text previews of HTML messages being two notable examples.

The calendar is an embedded version of *Korganizer* which is largely unchanged (even down to the 'Configure *Korganizer*' entry in the Settings menu) and works well.

The 'Home' page can display a selection of *rss* feeds and local (or not so local) weather, and both options come with a simple interface to add new instances. You can set an interval for updates but not much else, and if you've used *Evolution's* excellent Summary you'll be missing the extra details like upcoming appointments and an overview of your email activity.

Conclusion

In fact, the tools on offer as standard don't compare that well to the main competitors, but *Aethera's* unique selling point is the integration of TheKCompany's *Jabber* and *Whiteboard* plugins which elevate the product

greatly. The use of *Jabber* with its open standards and extensive range of clients and servers, makes this a good choice for the corporate network. As does the availability of both Windows and Linux versions. Both show promise and the basics are well covered, but, especially in the *Jabber* client, there are still a few omissions. Rendered emoticons, for instance, are by no means essential, but they add an element of fun and can, on occasion, help express 'mode of address' better than their plain text equivalents. TheKCompany says a new version of the *Jabber* plug in is imminent and this will provide file transfer tools.

The *Whiteboard* works well and contains some useful features especially if you're working on a Zaurus PDA. The tools could be more extensive, but it's probably unrealistic to expect a *GIMP*-like experience! Once over the [fiddly] setup procedure, use is easy and even quite responsive when communicating with users on a dial-up connection.

Overall it's not a bad package, but if you don't need the *Whiteboard*, you're probably better off with a combination of *Evolution* and *Kopete* (KDE's IM client) which can use MSN, AOL etc or a dedicated *Jabber* client. **LXF**

VERDICT

Features	5/10
Performance	6/10
Ease of use	8/10
Value for money (plugins)	6/10

An average product facing very strong competition for the corporate groupware crown.

LINUX FORMAT RATING
6/10

HARDWARE

Linux input devices

PART 1 Linux isn't all about servers and services – thanks to the continuing efforts of the community, it's turning into a first-class gaming and home use platform too. Paul Hudson highlights the hardware that works with Linux...

Keyboards, mice, joysticks, and steering wheels – the way we interact with our desktops continues to evolve as the drive for comfort, usability, and features pushes technology forward. For a long time, our favourite operating system steadfastly refused to support home users, which left cool new kit such as force-feedback steering wheels somewhat out in the cold. However, a lot of work has gone into the development of software for home peripherals, and we figured it was about

time we took a look at how the market stood – what hardware was on offer, how well it performs, and how much compatibility you can expect.

If you've ever per chased voice wreck ignition soft wear, you'll know it's not really up-to-scratch yet. As such, nearly all of us spend lots of time developing the love/hate relationship we have with our keyboards. While keyboards are increasingly seen as a neat place to drop cake crumbs, the reality is that they haven't changed much. Granted, you don't see rubber

keys any more (popularised by the 48K Sinclair Spectrum!), and the occasional keyboard is split in the middle to help reduce RSI, but the concept hasn't budged the width of a Caps Lock light. So, modern keyboards compete on the basis of looks, design, and 'feel', and we reviewed three of the most popular and advanced keyboards currently on sale.

The main problem with the keyboards we reviewed was getting the special keys to work properly, however even that just took ten minutes or so. While it would be nice if manufacturers

included Linux configuration support inside these products, that's not likely to happen for a while. In the meantime, tools like *xev* are essential, and you should check through your old copies of *LXF* if you experience problems, as they can be a goldmine of information.

Next issue, we'll be looking at force-feedback steering wheels and other gaming devices – things you wouldn't often associate with Linux. However, thanks to the ever-increasing popularity of Linux as a gaming OS, this sort of hardware is at last within reach.

Logitech Cordless Desktop Comfort

BUYER INFO

If RSI is an issue, this ergonomic Linux-compatible keyboard could be what your wrists are crying out for.

- **DEVELOPER** Logitech
- **PRICE** £99.99 (inc. VAT)
- **WEB** www.logitech.com

This Logitech natural-style keyboard and mouse combination pack retails for £100, which seems an awful lot at first considering that it's just a mouse and keyboard. However, the fact that they both work without cords gives you a lot more freedom on your desktop – as long as you don't mind changing the batteries once in a while, you'll find the product easy to use, easy to type on, and actually not bad for the price.

The Logitech radio receiver is wired to your PC using standard PS/2

connections, meaning that you plug them in as you would your keyboard and mouse – no need to try to get things working over USB. Getting all the keyboard's extra keys to do something useful was simply a matter of using *xev* to first find the key code they generate then mapping them onto something useful – I was unable to get the Sleep button to work correctly, but that's a problem I've seen previously in other keyboards. Hoyt Duff's tutorial on getting extra keys to work in Logitech keyboards (*LXF28*) worked perfectly here, which made life easy.

The actual working of the keyboard is very good – you get a

rather nice 'click' sound when a key has been pressed properly, which is something I always appreciate when programming. The mouse was fairly good, but the most interesting thing was the range of the wireless connector – we tested it sitting over two metres from the screen, and it all worked perfectly. As a result, we could set the screen fonts to a much larger size, sit further from the screen, and bid eye strain goodbye.

Overall, the Logitech combo is a great unit, made by a well-respected hardware company with a spotless reputation for quality and fine design. While the price might seem high at first, we encourage you to give the unit a try – you won't regret it!

Wireless performance was faultless – free yourself from cable clutter with the Logitech keyboard.



VERDICT

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

A great combo set, but not the cheapest price around.

LINUX FORMAT RATING

■■■■■■■■■■ 8/10

Microsoft Optical Desktop Pro

BUYER INFO

Surprised to see a Microsoft device in a Linux Magazine? *LXF* doesn't discriminate – it's a great product!

- **DEVELOPER** Microsoft
- **PRICE** £79.99 (inc. VAT)
- **WEB** www.microsoft.com

The mouse wheel really has a solid movement to it that gives an excellent feel of quality. Many Linux users hear "Microsoft" and run a mile, which is a shame really because MS have a long history of making quality hardware. This natural-shaped wireless mouse/keyboard combo is made entirely of very smooth plastic in two-tone silver and grey colours, giving a very metallic look but a very plastic feel – if you own a dark-coloured PC, such as most Dell or Alienware machines, this would look great alongside. The MS keyboard has the most 'special' keys in the test at

seventeen, with the Sleep key, cleverly, being sunken into the keyboard – no more will accidentally nudging your keyboard against your monitor turn your computer off mid-*Quake!*

Like the Logitech, the receiver plugs directly into PS/2 ports, meaning that configuring the mouse and keyboard was a doddle. Each of the hot keys worked perfectly and were easily configured. The mouse isn't quite as comfortable to hold as Logitech's, but the keyboard is much quieter, and the

mouse wheel really has a solid movement to it that gives an excellent feel of quality. One minor downside is that the receiver doesn't seem to have quite the same range as the Logitech equivalent, and sometimes we even experienced problems when a table was in between keyboard and receiver.

Compared to the Logitech keyboard, the MS offering does fairly well. It's lighter, has a more attractive design, and the special keys are larger. Having said that, the mouse is shaped like Microsoft's basic optical mouse, and the

Despite glitches with the receiver, this cordless setup beat the competition.



receiver needs a little tuning, however the disadvantages are more than outweighed by the advantages, and, at a price of £80, it's fair to say that the Microsoft offering has the edge.

VERDICT

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

Suffers from poor receiver, but otherwise leads this pack.

LINUX FORMAT RATING

9/10

IBM USB keyboard with UltraNav

BUYER INFO

Pushed for space, or want sleek, slimline looks to compliment your hi-tech surroundings? Choose this!

- **SUPPLIER** IBM
- **PRICE** \$99.99
- **WEB** www.ibm.com

This is a space-saving keyboard and mouse combination that only has a footprint of about 40cm. As the mouse is built-in (you have the option of a trackpad or the 'UltraNav' nipple-like navigation device) there's no need to move your hands away from the unit, which saves an enormous amount of desk space. The keyboard keys feel a little tacky because of the thinness of them – you feel like you're permanently using a laptop. Furthermore, the mouse buttons are dreadful and need a bit of rethinking.

We had some ups and downs getting this to work under Linux. Firstly, the special keys ('Access IBM', volume up, down, and mute) simply refused to

If you're used to using a laptop, you'll welcome the integral trackpad.

work, which was quite frustrating. The keyboard doesn't have nearly as many of these special keys as the other two here, so not being able to get them work isn't too much of a shame, however you do wonder why IBM would think their customers would want a dedicated key to get to IBM's website. The rest of the keyboard, which connects through just one USB port, worked perfectly as soon as it was plugged in, including two neat little keys next to the cursors which cycle through terminals – quite cool, required no configuration, and as addictively helpful as mouse wheels!

The nipple/trackpad didn't work to start with, but we didn't expect it to because X can be quite picky that way. However, after installing GPM and setting it up with default parameters, both the nipple and trackpad worked perfectly.

On the usage front, the IBM is clearly the thinnest, smallest, and lightest of the three keyboards reviewed here, and it does win marks for that – if visual appearance is important, then this is the looker.



VERDICT

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

Takes portability to the extreme, but at the expense of comfort.

LINUX FORMAT RATING

7/10

Algorithms for Compiler Design

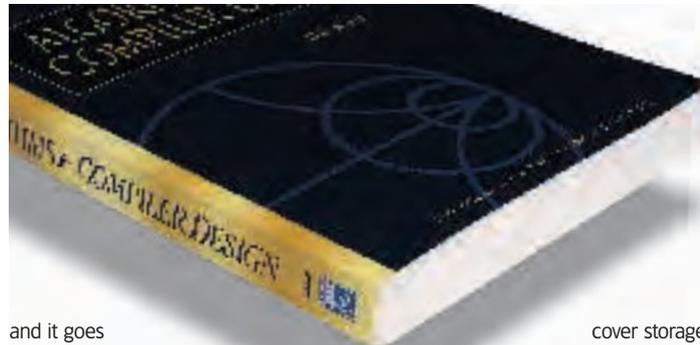
Paul Hudson presents a new book that might catch the interest of the compiler cognoscenti.

BUYER INFO

- **AUTHOR** O.G. Kakde
- **PUBLISHER** Charles River
- **ISBN** ISBN: 1-5845-0100-6
- **PRICE** £41.95 Pages: 334

There are several “greats” in the collection of books that cover compiler technology, and the high quality of these few have made the barrier to entry quite high for newcomers. If you’ve been following SKYLang from the outset and have yet to pick up a book to follow along, it’s perhaps time that you did. Unfortunately, this book, when compared to the greats in the area, doesn’t make the grade by quite a margin. However, that’s mostly because it isn’t particularly suited for novices, as takes a different tactic to the others; this is a book unashamedly devoted to algorithms.

Over half the book is there to cover the various ways to parse text,



and it goes into a great deal of depth about the various options. If you’re not a fan of flow charts and the like, this almost certainly won’t be for you – there are dozens of illustrations throughout that shed light upon the different techniques discussed.

Being such a small book, at least compared to its competitors, there is very little space dedicated to non-parsing matters – just five pages discuss symbol tables, eight pages

cover storage management, and nine pages discuss error handling. This is not nearly enough space to go into any sort of serious depth about these topics, and really should have been either covered properly, or left out completely rather than included in this truncated and confusing form.

The only sizeable, non-parsing chapter is that on optimisation, which gets 24 pages devoted to it. However, even with that the description is very

abstract, focusing more on “this is what a technique is”, as opposed to “this is how to implement a technique”. For some this isn’t a problem, as at the very least the book is language-agnostic. However, without the kind of concrete examples most people like to see, for reference if nothing else, you’re left almost entirely to your own devices.

If you want a particularly detailed discussion of parsers with only a very high-level discussion of other parts of a compiler, this book might be for you. However, at the price, there are many better options, most of which cover parsing in as much depth.

VERDICT

Not bad if you’re especially interested in parsing technology, but not worth the high price otherwise.

LINUX FORMAT RATING

6/10

Linux in a Nutshell

That NUT doesn’t, as Richard Cobbett discovers, stand for New User Training.

BUYER INFO

- **AUTHOR** Ellen Siever, Stephen Figgins and Aaron Weber
- **PUBLISHER** O’Reilly
- **ISBN** 0-596-00482-6
- **PRICE** £28.50 Pages 812

This is the king of the infodump. The Kaiser of the keyboard. The Emperor of the Enter button. And any other bad alliteration that should spring to mind. *Linux In A Nutshell* is easy enough to describe, but its sheer scale requires a handy thesaurus on hand, not to mention a packet of headache pills. It is, in short, a collection of every major Linux terminal command that you’re ever likely to need, all packaged up into alphabetical order, explained in plain English.

Criticism first. While undoubtedly useful, *Linux In A Nutshell* has the traditional dictionary problem that you have little to no chance of tracking



something down if you don’t know what it is or how to spell it. This is a guide where *sendmail* (a mail transfer agent), *sed* (stream editing) and *sdiff* (Finding the difference between two files) share a double page spread, and there are no convenient categories to help slim your sample set down from ‘the entire Linux command set’. While this will always be a problem, a handful more dedicated chapters to carve off individual applications, file functions and the like would make the inevitable scavenger hunts more

manageable. Provided that you do know what you’re looking for however, it’s hard to overstate the convenience of being able to refer to the book instead of constantly flicking between man screens and other guides. You get a full description of each control, along with both the short and long form of each attribute and a straight explanation of exactly what it does – be it obvious, like **-ignore-case** ignoring case changes, or **sendmail -am** using the *sendmail.cf* configuration. The final part of the

book takes a brief look at graphical desktops, although diving a little further under the surface than most starter guides simply by dint of pointing you to the actual directories where themes and other such configuration can be found instead of putting your faith in Wizards.

If you work in a larger Linux admin or development environment, you might do better to pay your \$14.99 subscription to O’Reilly Safari (<http://safari.oreilly.com>) and permanently add a copy of this to your list of ten browsing titles – if you don’t lock your office, this’ll be the first thing that a techie colleague will steal!

VERDICT

An essential reference for your desktop, but requires a certain level of knowledge for you to use it effectively.

LINUX FORMAT RATING

9/10

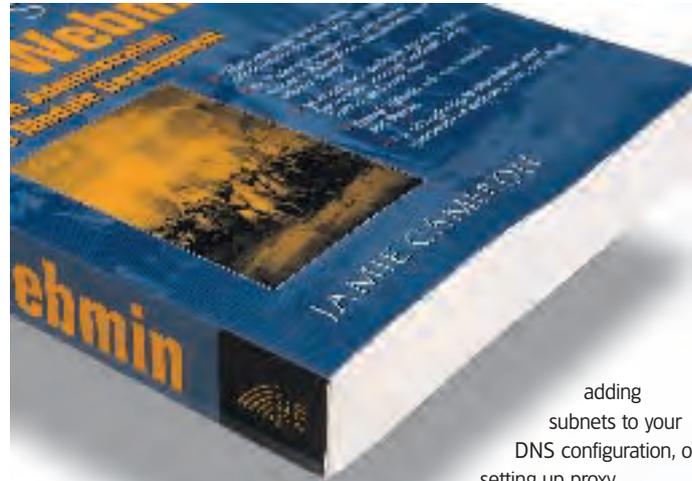
Managing Linux Systems with Webmin

Intuitively knowing when presence of mind means absence of body, **Nick Veitch** administers remotely...

BUYER INFO

- **AUTHOR** Jamie Cameron
- **PUBLISHER** Prentice Hall
- **ISBN** 0-13-140882-8
- **PRICE** \$44.99
- **PAGES** 794

Often when the authors of software sit down to write an explanatory book, the results can be disastrous. What seems like a good idea can turn bad in a number of ways. As is shown all too often in the Linux world (and the Open Source/Free Software sphere in general) what makes a good programmer or developer doesn't necessarily make a good writer of documentation. There is always the danger that authors will be too precious about his creation to impart any useful information; or indeed go into so much detail as to make the documentation worth little as purely a user-reference. It is a relief then that this book, while not exactly



entertaining, is at least frank, honest and accurate.

A lot of ground is covered in the book – in fact every official module is detailed here, with a thorough and unremitting description of all its features. Most of these are covered with goal oriented descriptions, eg

adding subnets to your DNS configuration, or setting up proxy authentication for *Squid*.

All of the sections are prefaced with a little explanation about the service the module is written for. While these probably wouldn't be enough for a complete beginner to get a total overview of things like mailservers and databases, but for some of the simpler modules there is just about enough info to get started.

The book is illustrated mainly with screenshots, which for the most part to be honest serve no useful purpose other to confirm that you are looking at the right page. At least most of them are reproduced large enough to see what is actually going on! There are also plenty of tables, and diagrams where necessary to help explain the concepts being discussed.

This isn't perhaps the most riveting read you have ever had, but it is clear, accurate and well explained, perhaps beyond the call of a simple explanation of the workings of *Webmin*. If you are considering the excellent *Webmin* for more than just home use, this book could probably save you a lot of trouble.

VERDICT

The best book on the subject. If you use *Webmin* extensively, you should simply buy it.

LINUX FORMAT RATING

9/10

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 "FUTLIN046".

	KEY FEATURES	DISK SPACE	MONTHLY TRAFFIC (Gb)	SHELL ACCESS	PHP4.3.2 SCRIPTING	MYSQL3.2.3 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	Yes	Yes	Yes	£8.95	£89.95	FREE	
LINUX GOLD	200	4	Yes	Yes	1	Yes	Yes	Yes	Yes	Yes	Yes	10	Yes	Yes	Yes	Yes	Yes	Yes	£14.95	£134.95	FREE	
LINUX GOLD PLUS	200	6	Yes	Yes	1	Yes	Yes	Yes	Yes	Yes	Yes	30	Yes	Yes	Yes	Yes	Yes	Yes	£19.95	£179.95	FREE	
LINUX PLATINUM	500	10	Yes	Yes	2	Yes	Yes	Yes	Yes	Yes	Yes	50	Yes	Yes	Yes	Yes	Yes	Yes	£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	Yes	£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

Operating Systems Fundamentals

If you fell asleep during your computing classes at school, Paul Hudson studies a catch-up.

BUYER INFO

- **AUTHOR** D. Irtegov
- **PUBLISHER** Charles River
- **ISBN** 1-58450-274-6
- **PRICE** £41.99
- **PAGES** 498

Charles River tends to publish 'above average' books – better than the bulk of books, but not the best out there. However, every once in a while the company comes up with a piece of magic that results in a particularly high-quality book that revolutionises the way people think about its topic. *Operating Systems Fundamentals* is not such a book. That's not to say it's a bad read – far from it; but it does take a certain kind of reader to fully enjoy what it offers.

Starting off with chapters such as *Introduction To Binary Arithmetic*, it moves swiftly onto data compression, then onto machine languages and assembly. The focus of the book is to



describe the workings of operating systems, which is a notoriously difficult thing to do – even more so considering this book has fewer than 500 pages.

However, the author has clearly worked hard at choosing the right blend of topics, and the book manages to cover each topic in a surprising amount of depth. How OSes load programs,

handle shared libraries and multitasking, how drivers and peripherals are supported, and much more – are all topics that are covered in a solid amount of depth, and make for quite a fascinating read.

What is noticeably lacking from the book, however, is specific application of techniques. The book, as it stands, thoroughly covers how components of operating systems work, and on that ground alone it does very well. However, it would benefit massively from more content, with discussion on OS-specific

implementation being most welcome. Ideally, the author would say "Now you've seen how virtual memory works, here's how Windows implements it, how Linux implements it, and how the two compare". While there are some mentions of application, for example how *ext2* handles inodes is covered, there really isn't enough of it.

Without enough of this practical text, the book remains almost entirely theoretical. This is no bad thing to a degree, but it does mean that the book will find most use in university and college courses, although readers who self-motivatedly study may also benefit.

VERDICT

Lots of very interesting content, but sorely lacking enough examples and detail in the practical department.

LINUX FORMAT RATING

8/10

Firewalls for Dummies

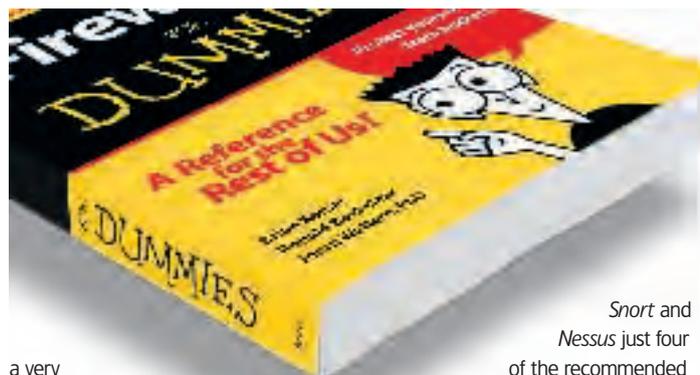
Richard Cobbett can't help but feel he's being told that his PC has already been protected...

BUYER INFO

- **AUTHOR** Brian Komar, Ronald Beekelaar, Joern Wettern
- **PUBLISHER** Dummies
- **ISBN** 0-7645-4048-3
- **PRICE** £17.50
- **PAGES** 428

Firewalls have never been more important. The growth of always-on connections, the sheer number of attacks currently making headlines around the world, and perhaps just as importantly, the increasing legal gray areas about your responsibilities as a PC owner mean that you cannot afford to go unprotected against hack attacks.

Simply installing Linux is often enough to provide sufficient security for a single PC, with almost every mainstream distro coming with a fully configured system, but mastering the wider techniques and terminology necessary to protect entire networks is



a very different story. You may not get much use out of *Firewalls For Dummies'* practical content, which devotes just one chapter to Linux and *iptables*, and then heads to Windows-land for a tour of the built in XP firewall, *ZoneAlarm*, *BlackICE* and co, but the theory elements are true for any system. In addition, the traditional *Dummies'* Part of Tens is surprisingly platform-agnostic, with *Netstat*, *Nmap*,

Short and Nessus just four of the recommended packages for your toolkit.

The theory section, as expected, requires little-to-no info – but would benefit from some stronger editing. Even some of the bread-and-butter elements, such as a brief overview of binary, comes across as far too garbled, even if explaining it should be as easy as 1, 01, 11. Luckily, the author soon finds his niche in the more involved elements, moving away from strained *Dummies'* 'humour' like "Whether you are an

engineer of any kind or not though, you probably realise that a computer firewall doesn't look like a brick wall" towards DMZs, CARP, and finally a surprisingly involved look at the Check Point Firewall-1. These technical sections don't offer an especially in-depth guide if you already know the basics of firewall admin, but they do point you in the right directions to continue your research. While the basic subject of advanced firewall configuration may not seem the ideal *Dummies'* fare, this book will at least get you over the initial technical learning curve to the point where you can start your practical planning. **LXF**

VERDICT

Neatly presented as always, but struggles to find a level between amateur and professional that caters for the Dummies market

LINUX FORMAT RATING

6/10

Roundup

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



Java IDEs

Richard Drummond puts three of the top free Java IDEs head-to-head.

Java tools range in price from zero monetary cost to several thousands of pounds. With such a great range in price, surely the IDEs at the low end of that spread cannot be any good? To find out, we are going to take a look at and compare three prominent, free Java IDEs: Borland's *JBuilder 9 Personal*, the Sun-sponsored *NetBeans* and the IBM-sponsored *Eclipse*.

As many readers of *Linux Format* will be aware, the word 'free' is overloaded in the English language. Here we mean free as in 'free beer' not 'free speech'. While *Eclipse* is entirely Open Source and is free in both senses of the word, *NetBeans* and *JBuilder* are not. *JBuilder Personal* is a proprietary product that is

available for download at no cost – but you don't get access to the source code and you can't redistribute it. *NetBeans*, on the other hand, is more confusing. The core *NetBeans Platform* is Open Source and hence doubly free like *Eclipse*. However, a useful version of *NetBeans*, which includes the necessary Java development tools, contains code distributed under a variety of licenses, both free and non-free.

So what can we expect in a free IDE? While these IDEs differ in the features they offer, they all provide rudimentary project management tools, an editor optimized for entering bug-free Java code (with goodies such as as syntax-highlighting, code-completion and navigations tools), and

facilities are running and debugging Java projects. They offer more, too, but in differing degrees. On the whole, programmers hate boring and repetitive tasks, and when programmers get bored, bugs creep in. Thankfully, computers excel at the dull and repetitive, so all of our IDEs include wizards and tools to automatically generate a lot of the boiler-plate code found in every programming project, to make sure it's bug-free.

All the whizz-bang tools in the world, won't make people use a product if it isn't pleasant and intuitive to use – and the three IDEs on test all tackle the same user interface problem in different ways – so we will have something to say about GUI design in what follows.

OUR SELECTION AT A GLANCE

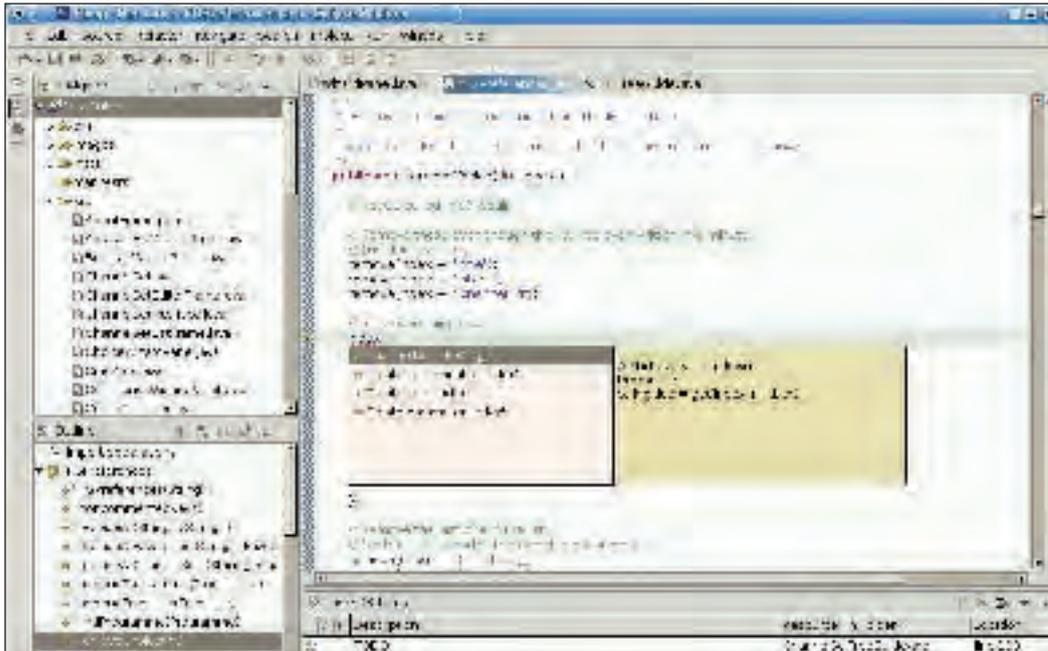
- Eclipse
- NetBeans
- JBuilder

'All three of the IDEs featured here include wizards and tools to automatically generate a lot of the code to make sure it's bug-free.'

Eclipse

IBM's new kid on the block is short on power tools, but would you take a look at that GUI!

■ VERSION 2.1.1 ■ WEB www.eclipse.org



Eclipse's SWT-based GUI is fast, responsive and provides a slick environment for editing Java code.

Eclipse is the youngest of the three IDEs on test, the project having been initiated when IBM open-sourced the core of its commercial Java development suite, *VisualAge for Java*, in November 2001. Thus *Eclipse* doesn't yet offer many of the heavyweight features of its two rivals: no GUI-building or bean building tools are provided and no tools for developing web-based or database-driven applications. However, these types of features and more are available in IBM's *Eclipse*-based replacement for *VisualAge*, now named *WebSphere Studio Application Developer*.

IBM is not the only vendor shipping products based on *Eclipse*. The purpose of the *Eclipse* project from the start was to provide an open tools platform with which a variety of development tools from different vendors could be integrated. For instance, Rational – famed for its UML design tools and recently purchased by IBM – also provides *Eclipse*-based tools and plugins. It's not just about Java development, either: C/C++ and Cobol tools are being developed as *Eclipse* sub-projects, and another effort is the Graphical Editor Framework Project,

which will eventually allow visual GUI-builder and UML tools to be implemented within *Eclipse*.

Although *Eclipse* may seem under-powered compared to *JBuilder 9 Personal* and *NetBeans*, it does provide a slick and attractive development environment which has some important differences from its two rivals. Most noticeably, *Eclipse* isn't built on Java's Swing GUI framework like the others. Rather, it employs its own *SWT* (*Standard Widget Toolkit*) which combines Swing's power with *AWT*'s native look and feel. Like the old *AWT*, *SWT* uses the host platform's native selection of components, *SWT* supplements these with pure Java components where necessary. On Linux, *SWT* can be built upon either *Motif* or (infinitely preferable in my view) *GTK+ 2.0*. With the latter, *SWT* applications look great, can employ a powerful array of widgets, integrate better with the rest of your desktop and respond much more crisply to user input than Swing apps tend to do.

The *Eclipse* workspace is called the Workbench, and it provides a fast and configurable place to do Java

development. An *Eclipse* window is typically split up into many panes which can be re-sized and re-arranged however you see fit.

Drag and drop

Drag-and-drop a pane on another's title-bar and it will be moved there as a tab-pane; drop a pane on another's left or right border and the target pane will be split horizontally, displaying both; drop a pane on another's top or bottom border and it will be split vertically. Each pane can hold a view or an editor. A view is tool for browsing and navigating a hierarchy of information. For example, the Resource View, shows the tree of folders and files in your current project, while the Package Explorer View gives Java's take on your project, showing a hierarchy of packages, classes, fields and methods, etc. Editors allow you modify some resource: for instance, the Java Editor provides a familiar editing environment will all the usual bells and whistles for writing and modifying Java code. *Eclipse* groups a collection of Views and Editors and their layout as a perspective. Default perspectives are defined to give easy access to the tools

you need for particular tasks. Thus there's the Resource Perspective, the Java Perspective, the Debugging Perspective, etc. You can customise these and create your own set of perspectives, if you need to.

A powerful aspect of the Java environment is the your project is rebuilt automatically every time you save changes to a file. This borrows a leaf from the *eXtreme Programming* book, and makes it easy to ensure that your projects are always kept in a buildable state. Errors are flagged in the editor and added top the to-do list of the built-in task manager. Hover the mouse over an erroneous line, and a suggested cause of the error will appear as a tool-tip. *Eclipse* even offers a Quick Fix function which will list possible solutions to an error and carry out the fix that you select. The Java editor is well stocked with code-completion tools, refactoring tools, and wizards to generate Javadoc comments, overload a parent class's methods, generate accessor methods and more. *Eclipse* offers poor support for viewing Javadoc documentation, and has no tools to aid the editing of HTML, JSP or XML pages – nor even a built HTML viewer. *Eclipse* spawns an external web-browser for viewing its own documentation (served up from a built-in web server) and external Javadoc pages.

Eclipse provides integration with *JUnit* for building test cases to exercise your software, and minimal support for *Ant*, the popular Java build system. Support for team-development with CVS is good.

VERDICT

Features	9/10
Ease of use	6/10
Documentation	8/10
Performance	7/10

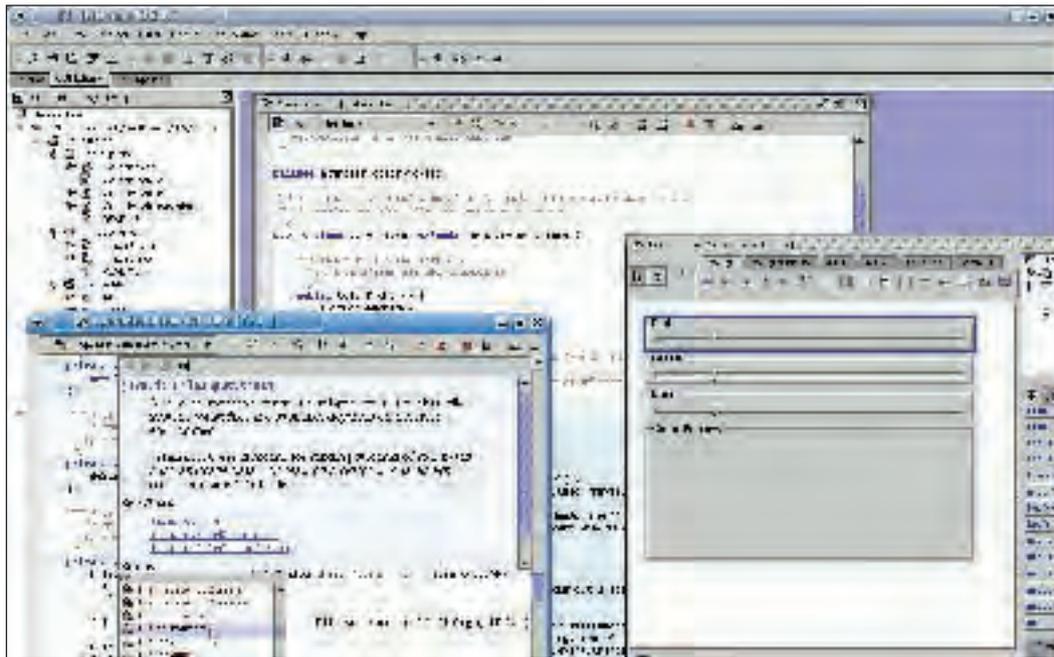
Eclipse currently lacks labour-saving tools for 'serious' development, but it has such a well-engineered interface that many will prefer over those of the more muscular *NetBeans* and *JBuilder*.

LINUX FORMAT RATING
7/10

NetBeans

The Open Source option from the creators of Java itself.

■ VERSION 3.5 ■ WEB www.netbeans.org



The *NetBeans's* underpowered source editor is balanced by great code-completion tools and a fantastic GUI builder.

Like *Eclipse*, *NetBeans* was derived from a commercial product. In this case, it was Sun's Forte for Java, but the birth of the *NetBeans* project predates *Eclipse* by some years. Sun's commercial offering that's based on *NetBeans* is now called *Sun ONE Studio* and is a full-blown J2EE environment complete with *Sun ONE Application Server* bundled with it. However, Sun's full-price product tends to trail the leading edge of *NetBeans* development, and *NetBeans* itself features plenty of J2EE tools.

Although nominally open-source, the licensing of *NetBeans* is actually rather complicated. The core of *NetBeans* – the *NetBeans Platform* – is Open Source, released under the Sun Public License, a variation of the *Mozilla* license. However, many of the tools which are necessary to turn *NetBeans* into a useful Java IDE are not open-source. A consequence of this is that binary downloads of the *NetBeans* IDE are only available under the Sun Binary Code License, which places restrictions on redistribution. Most users won't be bothered by such details, however, since *NetBeans* doesn't cost anything.

NetBeans has a flexible and configurable user interface. This works either in MDI mode – where all the *NetBeans* windows are embedded within a single main window – or SDI mode; where each *NetBeans* window is a top-level window in its own right. MDI mode is the default and makes the best use of screen real estate while still giving you a great deal of control. For example, you can float internal windows off as top-level windows, re-position internal windows as you wish, and re-attach a window as a tab-pane of another window. Like *Eclipse*, *NetBeans* lets you group windows in task-oriented groups, which *NetBeans* calls workspaces. To start with Editing, GUI Editing and Debugging workspaces are provided, but you can customise these quite easily and also add your own.

Explore the filesystem

The *NetBeans* Project Explorer Window gives you multiples views of the hierarchy of components in your projects. The Filesystem view shows a tree of files and any internal structure of those files (such as the classes, fields and methods in a Java source file). You must explicitly mount directories here for

NetBeans to make use of them, and this provides a lot of flexibility. You can also transparently mount and work with Zip archives or JARs here, or CVS repositories. The Project view shows a similar view of any projects you have defined. You don't have to add files to a project to edit and compile them and test the result, but this is an extra organisational tool for you to use. The third view in the Explorer is called Runtime, and this displays various runtime information, such as any Java processes you are running or debugging and any live database connections. You can also control the *NetBeans* internal web server here and, via the Server Registry, any other servers that *NetBeans* knows about. For example, *NetBeans* comes with a full install of the *Tomcat 4.0.6* servlet engine for testing web applications.

NetBeans is a modular IDE and the default build ships with modules for doing GUI development, building and using JavaBeans, working with XML and doing web-based development with servlets and JSP. Support for employing the Java build-system *Ant* also comes as default. Various other modules are available for download with the built-in

software updater, including support for RMI, CORBA and JNDI. Modules are also available to make use of testing frameworks such as *JUnit* and *Jemmy*. Wizards are supplied to take the mind-numbing spade-work out of generating common Java components, from plain classes and interfaces, to GUI forms, servlets and JSP pages. A JAR builder tool is also included.

With all these tools on offer, surprisingly one of the weaker areas of *NetBeans* is the source code editor. It provides a pleasant enough environment for editing Java, HTML, XML and JSP code, but it lacks configurability and is short on refactoring tools. There's no Quick Fix tools for diagnosing errors, and the code reformatting tool is laughably primitive. Having said that, the code-completion tools are excellent. Not only are possible completions listed in a pop-up, but any available documentation for the highlighted completion will be also be shown in a mini Javadoc viewer.

The *NetBeans* form editor is joy to use and provides a powerful tool for building GUIs using standard Java components and any JavaBeans you create. A particularly nice touch is that any code generated by the Form Editor will be highlighted in blue in a source file and protected from accidental editing (protected code is flagged by metatags that *NetBeans* embeds in code comments and are not displayed by the editor). A downside is that you cannot easily use the Form Editor to modify an externally create GUI class – *NetBeans* forms must have an accompanying XML file describing their structure.

VERDICT

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

Packed with loads of advanced tools and a great GUI editor, *NetBeans* is an excellent choice no matter what type of project you're tackling.

LINUX FORMAT RATING
 /////////////// 9/10

Java IDEs **THE VERDICT**

You can't get something for nothing, or so conventional wisdom would have us believe. Open Source development has spectacularly disproved this theory time and time again, and we have re-evaluated our expectations accordingly. We now want a good deal for nothing, and thankfully these three Java IDEs deliver, albeit in different ways.

The three products are free for different reasons too. *NetBeans* and *Eclipse* partly because Sun and IBM know the PR value of open-sourcing, partly because they want to leverage the power of open-source development, but mostly because they are both trying to build third party support for their development platforms – and the easiest way to do this is to make them open. Borland's *JBuilder Personal* is free, probably partly in response in to the free availability of the other two, and partly as a loss-leader to encourage the take-up of their full-price products.

Regular readers of the *Roundup* series in *LXF* will probably be lamenting the lack of table giving an at-a-glance comparison of features of the IDEs on test. That's because the sheer number of features provided by these products makes such a table rather worthless – have a look at *JBuilder's* 31-page list of features on the Borland website to see what I mean!

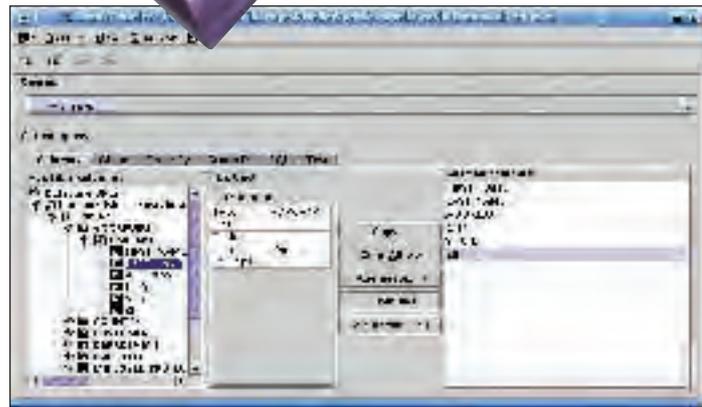
So, you're probably thinking now, which is the best IDE for me? It very much depends on the job in hand. If you're a hobbyist or a student or you just don't need the tools for web-based and database-driven development, then *Eclipse* provides such a slick editing environment that it would be hard not to recommend it. It has so many nice features that make you think "Hey, that's cool", that you can easily forgive it for its lack of 'real world' tools. On the other hand, if database development is your thing, then *JBuilder* is easily the best package of the three, and the only one that can claim to be "out-of-the-box" a rapid development environment for database tasks. For building GUIs and doing XML or web-based development, then the laurels go to *NetBeans*, which is also probably the best all round IDE of the three

IDE integration

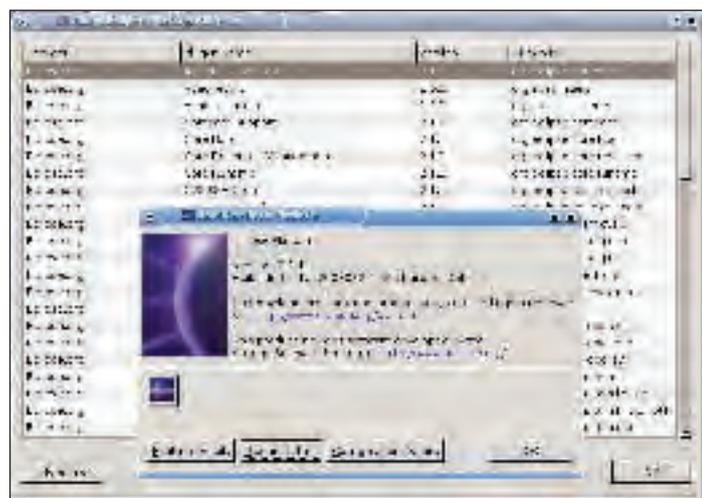
For those doing Java programming in the real world, then other factors will be prominent in your decision-making. How well do the IDEs in question integrate with the other tools you need to get your job done? In this regard, the strategies being employed by Sun and IBM are reaping rewards, since there's a lot more third-party development going on for the two Open Source platforms (the uptake of the *Eclipse* platform is particularly



The *JBuilder 9* boxed version is optimised for team-based development, with upgrade licenses from \$2,990. See *LXF44's* review where it earned a coveted *Top Stuff* award.



Pick the right tool for the job. *JBuilder* is well-equipped for DB development.



Both *NetBeans* and *Eclipse* have good third-party support for plug-in tools.

J2EE FOR FREE

Development on a budget

The *JBoss Application Server* (see www.jboss.org/) is an Open Source J2EE server, that can comfortably compete with expensive, commercial software such as BEA's *WebLogic* and IBM's *WebSphere*. One area in which *JBoss* is lacking, however, is GUI-based configuration tools. This problem can be solved by *Eclipse* or *NetBeans* with the appropriate plugin installed, and, hey presto, you have a full-featured environment for doing enterprise-level Java development – for no cost. Finding out more about these plugins can affect how you undertake your development project.

JBoss-IDE

www.jboss.org

JBoss-IDE is the official solution for providing *JBoss* with a GUI-driven development platform. This is a plug-in for *Eclipse*, and it sports tools for starting and stopping a *JBoss* server, and for packaging and deploying web applications.

NbJBoss

www.siptech.com/index?url=downloads

NbJBoss is a *NetBeans* plug-in which integrates the IDE with *JBoss*. It provides EJB templates, tools for packaging and deploying applications, support for managing *JBoss* servers and a GUI for editing *JBoss* configurations.

impressive given that it is the younger of the two). Do a search on sourceforge.net for *Eclipse* or *NetBeans* and you'll see what I mean. In the case of *Eclipse*, merely grabbing a selection of third-party plug-ins can do a lot to overcome its shortcoming for serious development (in particular, see *box J2EE For Free*). If you plan to use a particular vendor's J2EE server, then obviously you

should use the IDE that best supports that product.

As with a lot of things, though, the decision will come down ultimately to which IDE you prefer – which feels the best to you, which are you the most productive with. That's something that can't be answered until you have taken all three for a test drive. Since they are all free downloads, you've got nothing to lose except time. [LXF](#)

The new-look SUSE

SUSE is restructuring its efforts to maximise its offerings. Paul Hudson went to SUSE HQ in Germany and was given unrestricted access to its top echelons. Here's your chance to look inside SUSE, and see their past, present, and future...



Here's how the old SUSE chameleon looked...



And here's the new SUSE chameleon – only minor changes, but probably enough to ruffle some features amongst die-hards!

Having recently celebrated its tenth anniversary, few can dispute SUSE's long-term commitment to Linux. At a time when rumours are flying that Red Hat will withdraw its retail consumer offerings, we'd heard rumours that SUSE was planning to push its client packages further than ever, while at the same time expanding its server line-up.

Linux Format's Paul Hudson flew to SUSE's headquarters in Nuremberg, Germany, to find out what all the fuss was about, and spent two days talking to developers and managers alike to find out quite what was happening, and, more importantly, *why*.

Simply change

Despite being the one of the top distributions in the world, SUSE has no wish to rest on its laurels – when we visited, SUSE was already putting the

finishing touches to SUSE Linux 9.0, as well as an all-new product, SUSE Linux Standard Server. To top that off, SUSE has implemented a complete rebrand of the company and their products, with plans to bring all their products into line as soon as possible.

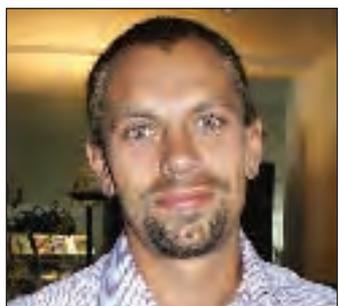
To facilitate a rebrand, SUSE's signature chameleon, "Geeko", has had a facelift, along with a new slogan, "Simply change", and product boxes have been redesigned to resemble each other more closely, to establish an enhanced coherent corporate identity. Perhaps the key change, however, is that SUSE has worked hard to produce a comprehensive roadmap for its products, in order to demonstrate its commitment to long-term support for enterprise users. In order to support this

new growth, SUSE is working to open up distribution and support channels in various countries around the world.

That's not to say consumers will be left out – far from it! While we were there, we got exclusive access to several key developers and managers behind SUSE's consumer product, and were very impressed at the amount of planning and innovation exhibited.

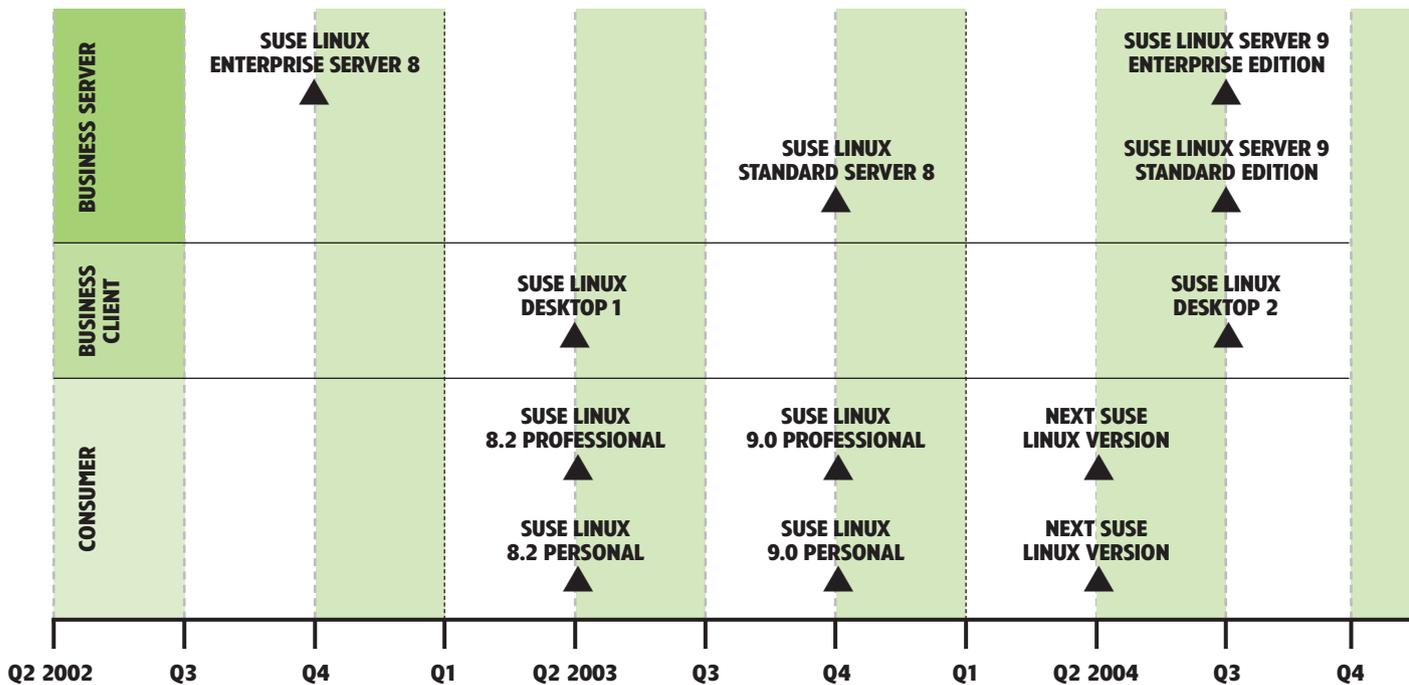
The new SUSE

With SLES raking praise in from all areas, Standard Server on the horizon, and SUSE Linux Desktop (SLD) already making major inroads into businesses everywhere, you'd be forgiven for thinking that SUSE had forgotten about its end-user product. That, however, is not so, as Dr Martin Sommer, the product manager for SUSE Linux Personal and Professional, was eager to point out.



"We strongly focus on the extension and qualification of our international channel network and will give our partners the knowledge they need and the skills they require to fully meet the needs of our mutual customers."

Thomas Zacharias Director International Partner Sales



While a lot of SUSE's effort appears to have been specifically geared towards the higher-end market, dozens of developers have been beavering away on the new SUSE release, SUSE Linux 9.0. If nothing else, this latest iteration of its product includes a slew of almost-invisible fixes that will improve the general feel of the system. For example, extensive work has been done to streamline boot times, and SUSE predicts many will see their computers boot up to twice as fast when using SUSE 9.0. Furthermore, more backports from the 2.6 test kernel have been done, giving users a tweaked multimedia environment.

On the topic of 2.6, SUSE is hoping to include a current snapshot build of the 2.6 kernel (pre-release, of course) inside 9.0 as an optional package for more advanced users. While it's not likely that many will take advantage of this kernel, it does at least give power users the chance to toy around with the latest and greatest software, as well as test out their own code against 2.6.

Going beyond the small features, there are of course the latest barrage of programs available. The big name is *OpenOffice.org 1.1*, which in itself contains hundreds of excellent new features such as improved *MS Office* compatibility, PDF export of documents, and a macro recorder, not to mention the startup

speed tweaks. SUSE 9.0 will also include *Scribus 1.0*, the increasingly popular new DTP package that we reviewed in *LXF45*.

There are also a variety of new features available in the latest release of *YaST*. Firstly, there have been a variety of tweaks to the networking setup – DSL modems are now detected and configured by default, as are up to 50% of all Winmodems.

There's now also easy-to-grasp support for DNS and DHCP for home network use. Perhaps the most interesting new feature in *YaST* is called the SUSE System Doctor, which works as a *YaST* repair module to reset the system to an earlier state in case of system failure by scanning for corrupt files and replacing as necessary. Naturally no one ever wants to be in the situation as to have to use this feature, but it's good to know that SUSE have planned ahead just in case.

We spoke to Dr Martin Sommer, the product manager of SUSE Linux Personal and Professional, about his plans for the distributions...

Linux Format: How important to SUSE is its continued support of home users?

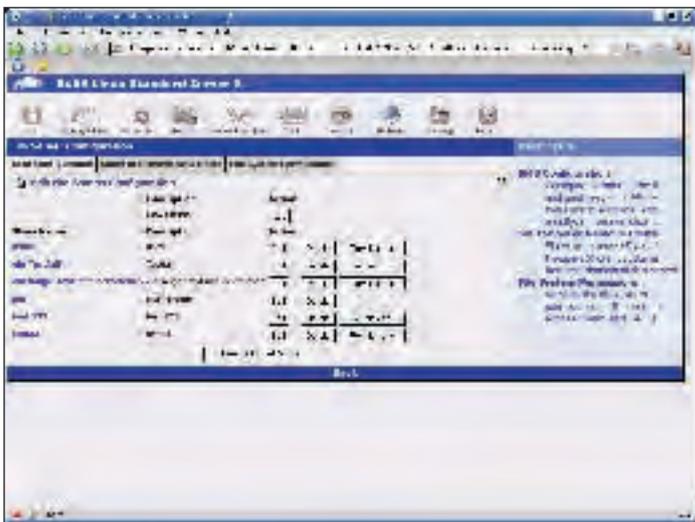
Martin Sommer: We have put much emphasis on our products for home users in the past and will continue to do so in the future. Our product for home users, SUSE Linux, is Europe's

The SUSE roadmap stretches deep into 2004, enabling its customers and partners to better plan their deployment schedules.



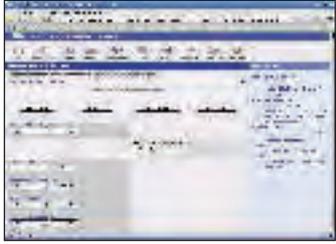
“SUSE Linux 9.0 combines desktop-comfort with feature-rich applications and brings state-of-the-art Linux technology to home PCs.”

Martin Sommer
Product Manager SUSE Linux Personal and Professional



With Windows NT 4 deployments being a key target market for SLSS, a lot of work has been done to make SMB sharing as easy as possible.

NEW-LOOK SUSE



Setting up your own DNS server is now just a wizard away in SLSS.



“SUSE Linux Standard Server’s detailed users access permissions, central user management with directory services and central server administration through web interfaces enable small businesses to maximise their productivity.”

Marcus Kraft
SUSE Product Manager

best-selling Linux package for home PCs and provides a wide variety of the latest Linux tools and applications, integrated in a feature-rich and easy-to-use desktop.

LXF: How do you think SUSE 9.0 stacks up against other home distros like Mandrake or Red Hat?

MS: Our customers especially benefit from the ease of use of SUSE Linux 9.0. We are the only Linux provider who offers a system-wide installation and configuration tool, *YaST2*, for both installation and all important settings including software installation. With *YaST2*, users can easily install SUSE Linux alongside an existing Windows operating system and can make use of SUSE Linux’s perfect compatibility with Windows systems. Moreover, technically experienced home users will appreciate the large number of ultimate Linux tools included in SUSE Linux 9.0. And not to forget, the depth of documentation SUSE Linux provides is enormous.

LXF: Which feature are you most proud of in SUSE 9.0?

MS: Definitely *YaST2* – with all its functionalities and its great modules is outstanding!

LXF: What has SUSE done to ease installation for first-time users?

MS: The biggest issue is that in most cases the user just selects his language and the rest of the installation is completely automated. In the ongoing installation process, users just need to click ‘OK’ or ‘Next’ and the SUSE Linux system is installed on their PC without any fuss.

Standard Server

While SUSE Linux Enterprise Server has proven to be an unprecedented success, thanks largely to its easy administration and smooth, cross-platform nature, many have found that it provides a little *too much* functionality for their needs. For example, small- and medium-sized companies have little need for load balancing, so why should they pay for it?

To solve this problem SUSE has engaged in a lengthy customer review and feedback process to make sure they could accurately gauge what SME customers felt they wanted most from SUSE. The result of this process, after thousands of man hours of development, is SUSE Linux Standard Server, the baby brother of SLES. Built

on the same codebase, and named version 8 to reflect this, SLSS shares much of the same functionality as SLES, although much of the more powerful technology has been taken out to make it easier to grasp.

The key differences between SLES and SLSS are that SLSS will only be available for x86, running on a maximum of two CPUs, and it will not include technology such as load balancing, clustering, and high availability – systems that really don’t see much use outside of very large corporations. In addition, SLSS will be debuting a new GUI interface for administrators, which, running through a web browser, allows system admins to control large chunks of their server through a very straightforward interface.

Clearly a lot of work has gone into this – if you’re thinking “this sounds like *Webmin*,” then you’re thinking along the right lines, but in reality what SUSE has created is much more than that. As it’s fully-integrated with your workflow, as well as being designed to offer lots of control over the box, the new administration system works like a charm – the interface is very easy to pick up, and has intelligently placed options to make sure you don’t need to navigate around more than necessary. SUSE plans for this new GUI interface to be rolled out to SLES 8 machines sometime later this year as part of the maintenance contract, which will be a welcome addition for all SLES administrators.

Since *Samba* support is integrated so deeply into Standard Server, interoperability with Windows is clearly going to be very easy. As Windows NT edges ever closer towards the end of its shelf life, more and more organisations are faced with a stark decision: make the jump to Windows 2000/2003, or consider the alternatives. With finely grained access control lists, full domain controller support, and *Samba* shares configurable through the easy web interface, it’s clear that Standard Server is looking to be that alternative.

On the pricing front, Standard Server is pitched at half the price of Enterprise Server, although that’s really not going to matter to most people. Of key importance is the comprehensive support that SUSE provides – it offers top-to-bottom technical help for all

areas of server operation for a very reasonable price. For standard-level support, that is 9am–6pm with a solution turn around time of four hours, you pay only 300 Euros a year. Customers wanting even more from their maintenance contracts can pay 525 Euros a year for 24x7 support, with a turn around time of just two hours – a real bargain, and almost unheard of in today’s environment.

We spoke to Frank Lemser, the product manager for SLSS – here’s what he had to say...

Linux Format: How does Standard Server fit into your product line?

Frank Lemser: By expanding the SUSE Linux Server family with SUSE Linux Standard Server, we meet the requirements of small organisations for an easy-to-configure server operating system that provides all relevant server services for x86 hardware. Enterprises with complex, heterogeneous IT infrastructures benefit from the strengths of SUSE Linux Enterprise Server – efficient administration in large networks, performance in mission-critical areas, its enterprise features and its availability across all relevant hardware platforms.

LXF: How much of Standard Server was based on input from end users?

FL: Small organisations are looking for IT solutions that are easy to install and to configure. Hence, we have put much emphasis on the usability of SUSE Linux Standard Server. Its server services can be easily configured through graphical configuration wizards, e.g. the deployment as a Windows domain controller, as file and print server in Windows environments or as an Internet gateway and name server (DNS).

LXF: In your opinion, what’s the target market for Standard Server?

FL: The SUSE Linux Standard Server is tailored for the needs of small organisations and departments, who are looking for a reliable, future-proof, high-quality server solution. Often, employees of small companies administer the IT infrastructure in addition to their regular job. With SUSE Linux Standard Server, we provide them a server operating system that combines the advantages of Linux with a user-friendly installation and configuration process

and a reliable maintenance offering that ensures smooth, secure and cost-efficient server operations.

Especially small organisations, who have never used Linux before, will benefit from the ease of installation and configuration of SUSE Linux Standard Server.

LXF: Which feature of SLSS do you think will be most popular?

FL: Customers will certainly be excited about the usability of SUSE Linux Standard Server. With the help of the configuration wizards and a complete graphical user interface, a ready-to-use server OS can be implemented easily, saving both time and costs.

Simply change

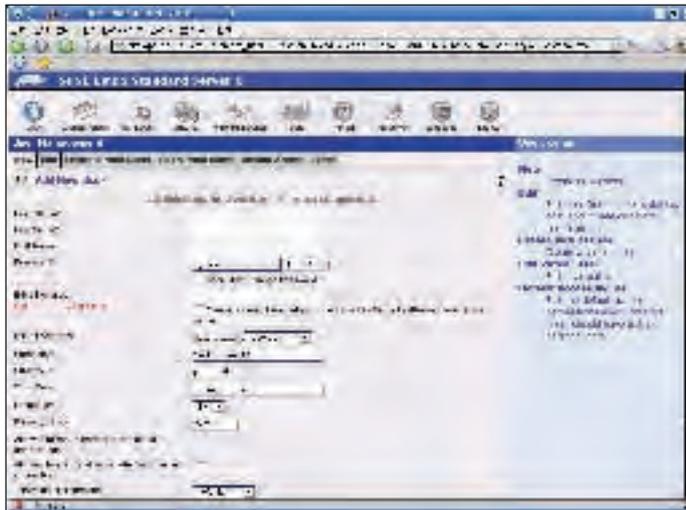
SUSE were founded on the principles of being the Linux experts – dropping the slogan doesn't mean they are no longer *The Linux Experts*, but instead shows them as being a progressive company, and one that's still working hard to innovate in the distro arena. With Red Hat's Enterprise Linux AS product scraping the verdict "Nothing too exciting" when it was reviewed in LXF42, it shows that SLES, and now SLSS, have got a lot of room to expand.

Despite all the hard work that has gone into SLSS, it's still quite possible that Windows NT admins will be less than eager to "simply change" away from their proprietary solution. However, the impending termination of Windows NT's support, combined with the excellent support program offered by SUSE, could well be enough to encourage droves of sysadmins to make the switch. With the excellent Samba support in SLSS, migration from Windows should be a fairly smooth process, ably helped by SUSE's dedicated support technicians.

What will be particularly interesting is the next release of the SUSE server family, due sometime in around Q3 2004, as this will bring all three



Adding DNS records is a point-and-click operation – as it should be.



Standard Server has finely grained access control, so adding a user lets you be as detailed as you want all within the same friendly interface.

corporate products (SLES, SLSS, and SLD) into line based around the 2.6 kernel, presenting a common front for businesses. Customers on SUSE's maintenance program will receive free upgrades to the latest versions as they are released, ensuring that they keep up-to-date with the newest software if they wish.

When we spoke to Juergen Geck, the Chief Technology Officer at SUSE, he said told us that "SUSE would rather not use version numbers for our enterprise products, because they aren't about having the latest and greatest software available. Instead, we want to provide support services, to help our customers use our software with the minimum of work".

It should be clear to all that SUSE's rebranding is more than just about tweaking the gecko and polishing its existing offering. Yes, SUSE has produced a comprehensive brand manual full of colours, shapes, and slogans, and yes, it has newly designed boxes that are more dynamic and cohesive to the larger theme of SUSE products, but all that is really just fluff compared to what has actually changed.

The key is that SUSE is making it crystal clear that its enterprise promise of supporting products for five years isn't just a publicity stunt – their roadmap makes it clear that they intend to keep their product line going well into next year, with much more besides. By firmly tying their products down to a predictable pace, partners such as IBM and HP are better able to work in tandem with

SUSE to make sure their customers get the most stable, reliable operating system possible. This was underlined as we toured through the SUSE testing suite, where dozens of machines from old Pentiums to high-end rack-mounted servers buzzed away 24 hours a day to make sure that "Certified for SUSE Linux" means a 100% guarantee of software and hardware compatibility.

21st century geeko

Today SUSE rely on the community as much as ever, although thanks to their resources – 150 full-time developers at the time of writing – the company is able to give back to the community lots of the code they produce. Being amongst the oldest of Linux companies, SUSE certainly isn't wont to taking risks; this new brand strategy is well thought through, considered, reconsidered, and planned until every last detail was worked out.

Overall, it seems like a good move for SUSE, as Red Hat has managed to at least give the appearance of a successful, enterprise-oriented business, even if its actual solutions lag behind a little – at the time of writing, there is still no Opteron support from Red Hat, and its consumer line is under threat of cancellation. With Mandrake still struggling to get its financial feet firmly back in the black, the way seems open for SUSE to make some serious in-roads, both in the consumer and enterprise market places. And, in our opinion, SUSE deserves this opportunity. **LXF**

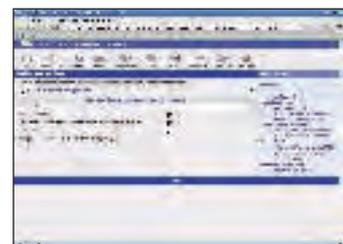
MORE INFORMATION

If you'd like to read more about SUSE business plans and service offerings, be sure to check out the website that is dedicated to just that end: <http://www.suse.com/en/business/products>



"I was very proud of our AMD64 port of Linux; it's now in the 2.4 kernel, which benefits everybody."

Chris Schlaeger
Director of Distribution Development, SUSE



Perhaps the most important tool is the one to help configure email.

HotPicks

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 also provide 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 or Mike Saunders himself at mike@aster.fsnet.co.uk

HotPicks at a glance

Kroupware	37
Danpei	38
gPHPEdit	38
MozPlayerXP	39
Straw	39
Cherokee	40
ScummVM	40
Atlantik	41
GRhino	41

HOTPICKS AWARD

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

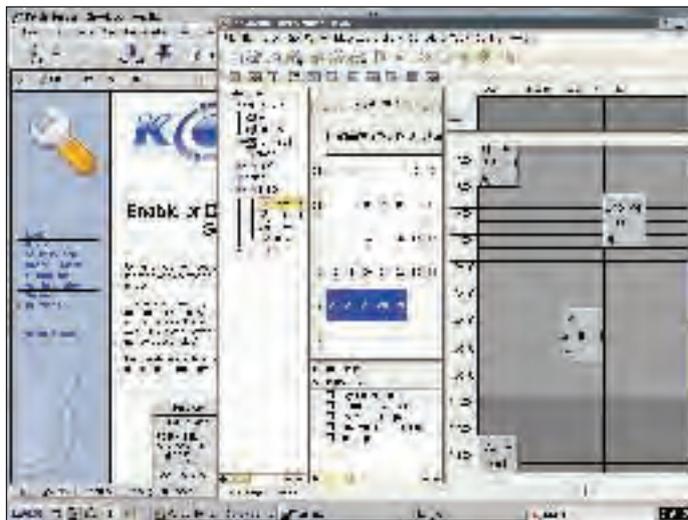


GROUPWARE SERVER/CLIENT

Kroupware



■ VERSION 1.0 ■ WEB www.kroupware.org



Kroupware's Kolabified KMail beside the web-based admin interface.

Small companies and organisations

usually find that basic email facilities are adequate enough for maintaining good communication amongst employees. In larger and more complex scenarios, though, a more comprehensive solution is required and several 'groupware' suites exist to provide this. Typically, a groupware system goes far beyond mere email functionality – it provides calendaring (for organising meetings etc), task lists, real-time discussion, contact databases and more, all shared through a central server.

Open Source tools in this vein for Linux include *phpGroupWare* and *OpenGroupware.org*; we reviewed the latter last month, and found despite showing a lot of potential, its lack of native clients and unintuitive Web-based front-end to limit its appeal. Kroupware, like *OGoo*, also expands on a handful of free software projects, but the main difference is that its KDE-based client (built around modified *KMail* and *KOrganizer* components) is ready for use now.

The German government and three companies (Erfrakon, Intevation and Klarälvdalens Datakonsult) sponsored the *Kroupware* project to provide a fully open alternative to the *Microsoft Outlook* and *Exchange* system, with the server and client being named *Kolab*. Installation can be tricky, particularly when building the server from source; the developers have opted to use the *OpenPKG* system which provides a uniform set of supporting packages across all distros.

Because *Kroupware's* server uses *Apache*, *Postfix* and other well-known progs, *OpenPKG's* guarantee of specific versions is a great help to the developers. The end result is a self-contained server suite which sits in */kolab* and doesn't interfere with the rest of the distro. Smart move! Installing from binaries is easiest – the docs give a quick walkthrough, and we've included the necessary files to get it running on Debian Woody (3.0). For the client, you'll need KDE 3.1 and no *kdenetwork* or *kdepim* packages installed – it's built

from modified versions of those additions, so just compile as usual.

Up and running

On the server side, *Kolab* makes use of SASL and OpenLDAP for managing authentication and configuration details. The email components consist of *Cyrus IMAP* to provide IMAP and POP3 access, along with *Postfix* as the MTA. In all, it's a good base made up from known, solid software. Admins will find the PHP-based web configuration front-end – which assists in enabling various services and managing user accounts – to be clean and hassle-free.

Regular KDE users will have few problems getting to grips with the KDE client, as it's essentially made up of *KMail*, *KOrganizer* etc with some extras sprinkled on top. *KMail's* handbook has been updated to reflect the changes – the developers have done a praiseworthy job. Alongside the core email functionality, shared contacts and task lists are available, as is *KNotes* for quick snippets of text – and all this is very straightforward to get to grips with.

Kolab's calendaring software (the modified *KOrganizer*) allows events to be planned and shared, giving the option of making available your 'free time' information, along with a to-do list for outstanding tasks. Without doubt it isn't the most exhaustive solution ever, but nevertheless does its job well. Sadly, there's no free and open source *Outlook* connectivity option available at present (Bynari offers a proprietary plugin) but <http://otilkcon.sf.net> may deliver an alternative in the future.

All things considered, *Kroupware* is at a similar stage to the *Mozilla* and *OpenOffice.org* 1.0 releases – good enough for day-to-day use, but awaiting the refinement that a wider user and developer base will bring. It's naturally unappealing to companies that already use GNOME desktops, and maintaining separate *kdenetwork* and *kdepim* packages could be a chore, but these issues are due to be ironed out in future KDE releases (*Kontact* is in the pipeline for even tighter integration). *Kroupware* is doing a fine job improving the validity of free UNIX desktops in the workplace.

IMAGE BROWSER

Danpei

■ VERSION 2.9.2 ■ WEB <http://danpei.sourceforge.net>

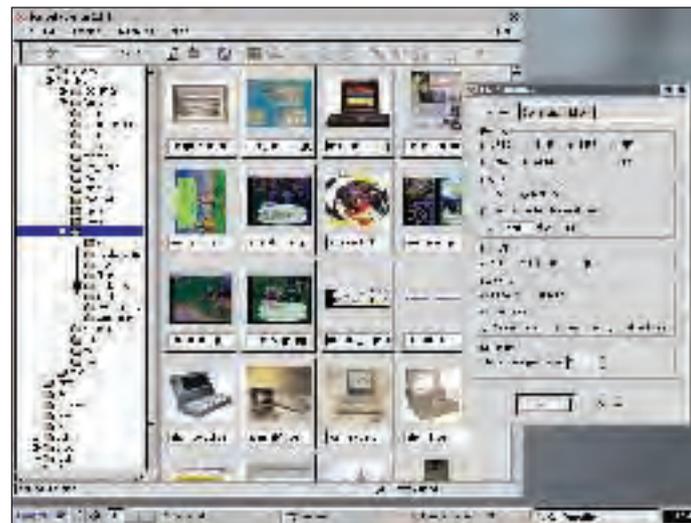
As digital cameras and good quality scanners continue to plummet in price, and the Internet grows ever more popular, increasing numbers of computer users nowadays have large collections of images stashed on their machines. If you're managing website graphics or a photo album, the traditional bundle of a generic file manager and viewer like *xloadimage* just doesn't cut the mustard any more – specific image management software is needed for browsing and manipulating your pictures. Richard Drummond looked at the KDE-based *PixiePlus* in LXF38's *HotPicks*, and now it's the turn of *Danpei*.

Being built around GTK, you'll need version 1.2.6 (or newer) of the toolkit to compile *Danpei* from source. The only other requirements are *ImageMagick*, *gdk-pixbuf* and *libpng*; these are installed on virtually every

major distro so you should have no problems with the build process. The resulting binary is pleasingly tiny at only 165KB (stripped).

Danpei's interface is a very approachable *Windows Explorer*-like affair – so non-Linux users will comfortably be able to browse your photos unsupervised. There's a tree pane down the left for navigating the filesystem and a thumbnail display on the right. Images can be sorted depending on their name, date, access and modification time, while three sizes of thumbnail size can be selected. It's a breeze to use and – most importantly – it's ultra-fast.

Currently, the program will read JPG, PNG, GIF, PCX, BMP, XPM, XBM and TIFF formats, and it uses 'display' from *ImageMagick* for full-size views and *The GIMP* for editing. These aren't hard coded though – they can be



Danpei in action, with the browser and Configuration box open.

altered in the brief settings box, along with confirmation dialogs and the thumbnail per page generation amount (100 by default for speed) among others.

Danpei doesn't intend to be an image editing tool and consequently the only operation that can be performed is rotation. Still, if the developers can add more hooks to the

ImageMagick functions there's plenty worth including in future releases, and although the occasional 'Japlish' text needs cleaning up it's not crying out for documentation. *Danpei* isn't as fully-fledged as *PixiePlus* by any means, but for quick image browsing it works without giggles. Zippy, easy and solid – a useful light image-centric alternative to *Nautilus* and *Konqueror*.

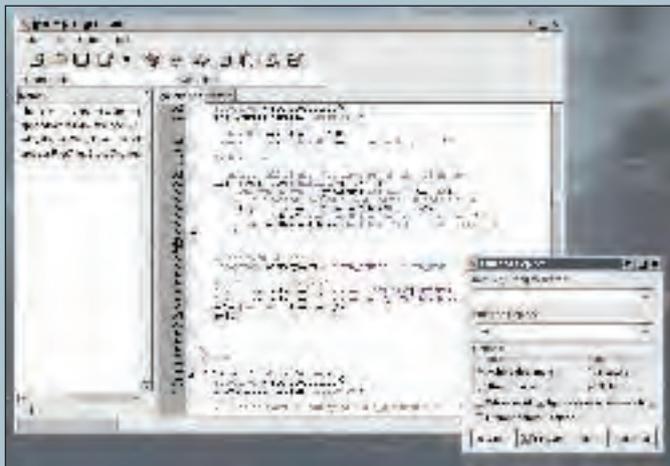
PHP CODE EDITOR

gPHPEdit

■ VERSION 0.4.3 ■ WEB www.gphpedit.org

PHP, formerly the 'Personal Home Page' tools and now the 'PHP Hypertext Preprocessor' (celebrating the hacker tradition of

recursive acronyms) is a scripting language for the Web with a gigantic following. It's also ideal for other general-purpose coding chores, and if



PHPers will enjoy smart syntax highlighting and quick navigation features.

you've been following Paul Hudson's PHP tutorials in recent issues of *Linux Format*, you may be looking for a specialised code editor. Programs for hacking C, HTML and other languages have existed for some time; now PHP is starting to see its own dedicated software and *gPHPEdit* is looking notably strong here.

After we put *gPHPEdit* on our coverdisc a few issues back, lead developer Andy Jeffries asked if we could give it some coverage in *Hot Picks*. It's been a while since the last release – earlier this year – but the editor already sports a nifty sum of features with a good deal more in the works. Being built around GNOME 2, you'll need the GTK2 and GNOME development packages installed if you're building from source. Similarly, the *Scintilla* editing component is required, and we've supplied RPMs to make things quicker. It should be noted that *gPHPEdit* has not had a stable release yet, but many of the common problems that people encounter when using it have been ironed out – see the bug tracker on the website for more details.

gPHPEdit's clean and unintrusive interface follows the usual window furniture strategy of a two-pane display, and toolbar for access to common operations. The left pane contains a class list for speedy navigation around large files, while the right holds the main syntax-highlighted editing box. Satisfyingly, *gPHPEdit* makes use of tabs for multiple open documents, and the editor allows code blocks to be collapsed and expanded for readability. Well crafted.

Another handy addition is the ability to check the current code for errors via the php binary itself, as is the drop-down list for long function names and the macro recorder (although it'd be a bonus if it could store more than one). The *Scintilla* component performs excellently; it's fast and doesn't suffer from the odd cosmetic glitch as we've seen with others. *gPHPEdit* does work commendably well as a specialised editor without getting bogged-down in IDE functionality, and is swift and reliable to boot. All regular PHPers should give it a good test run.

MOZILLA MPLAYER PLUGIN

MozPlayerXP

■ VERSION N/A ■ WEB <http://nexflo.net/users/mozplayerxp/>



At last, QuickTime trailers can be viewed straight in the Mozilla window.

Linux users who have been following the OS's progress over the last few years will be aware of the major problem in gaining desktop

popularity: integration. While we now have a broad range of top-notch Open Source applications and two very strong desktop environments,

fitting everything together as a cohesive whole is a massive task (although freedesktop.org and Red Hat's Bluecurve are helping here). This is just the downside of massively distributed development; everyone has different ideas on how things should be done.

Mozilla and MPlayer are perfect examples of this – they're both superb projects, but also lacking any integration. Thankfully, *MozPlayerXP* has arrived to let Mozilla users view MPlayer's video formats in the browser window. As a work-in-progress, the site holds little information on installing the plugin – for most users, simply extracting the archive in the Mozilla installation plugin tree will suffice (usually contained in `/usr/lib/mozilla/plugins` or similar). Note that you'll have to run `bunzip2` on the plugin archive before `tar xfzv`, as the developers have compressed it twice.

Building from source is the other option if the plugin fails, but this is considerably more time-consuming as it requires a whole Mozilla source tree to work with (typically eating up a few hundred MB of precious hard disk space). Ideally, *MozPlayerXP*'s coders

will accept ready made pre-built plugins for various distros and their `gcc/libstdc++/Mozilla` version combos. Lastly, of course, you'll need MPlayer itself – if downloading doesn't appeal, check out issue 44's coverdisc for the full version and LXF42 *HotPicks* for a review of its features.

With everything in place, Mozilla should now incorporate an MPlayer window whenever you hit a page containing the relevant media. Depending on your configuration, MPlayer will handle WMV, ASF, QuickTime, RealMedia and stacks more; it pretty much covers all the usual formats you're likely to find on regular sites. Simple it may be, but it works well – nonetheless, the coders behind it have a rewrite in the pipeline to fix a few known bugs. Despite there also being a few issues with Java calls that users may encounter, and the odd instance of pictures not loading in a window properly, if you're a Mozilla fan and view a lot of video-laden sites, give it a go. *MozPlayerXP* will also work with other Mozilla-based browsers, such as Netscape, Galeon, and Konqueror. Providing feedback on the forums that are associated with *MozPlayerXP* helps the developers.

NEWS FEED AGGREGATOR

Straw

■ VERSION 0.19 ■ WEB www.nongnu.org/straw/

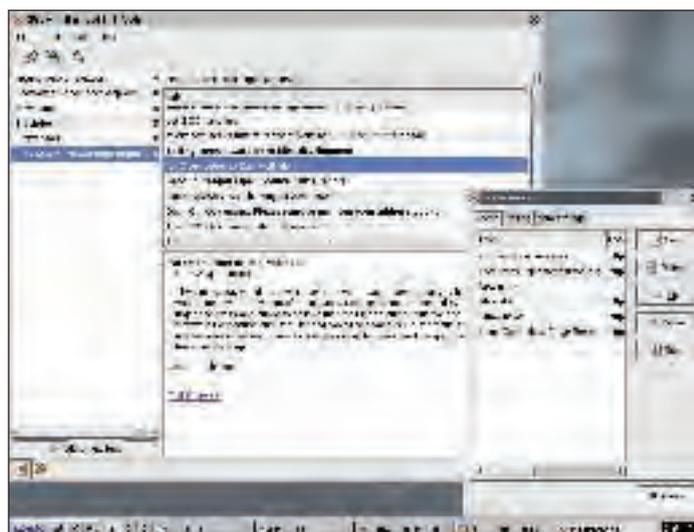
Web browsers have been built up so much, that they're now the catch-all solution to most online information. But for many simple jobs, such as grabbing the latest quick news snippets from a variety of sources, they're somewhat overwhelming in the features they provide – a lighter alternative would be better. Many site authors have recognised this, and have started offering 'RSS' (Really Simple Syndication or RDF Site Summary) feeds which deliver just the bare info to a client. This has been popular on news sites and weblogs, and *Straw* is a GNOME 2 client for RSS 'desktop news aggregation'.

Although *Straw*'s GNOME requirements are straightforward, it relies on a large array of Python tools and libraries to run correctly. Along with Python 2.2, ADNS, PyADNS,

`mxDateTime` and `PyXML`, the commonplace `PyGTK` and `GNOME-Python` bindings are required. Many of these are installed by the mainstream Linux distributions, but if you encounter difficulties, we've supplied all the extras you need on our coverdisc. It's always awkward when a program depends on so many seemingly esoteric bits and bobs and maybe the developers can fold in a few.

Aside from the lack of tooltips for *Straw*'s main toolbar, the interface is a cinch to operate and is built up from a site list down the left and article viewing pane on the right. The pane dimensions can be resized to make maximum use of screen space, and on the whole it's simple and workable.

When adding a new feed, you're prompted for a direct link to the RSS feed location, or simply the base URL



Looks like SCO is still 'clutching at Straws! Geddit?!

of a site, leaving *Straw* to find it automatically. Additionally, the polling duration can be altered from its default of 20 minutes – saving the need for manual refreshes. Having an audible alert for new items would be good to see, though.

Straw's aim is to be a 'faster, easier and more accessible way to read news and blogs than the traditional browser,'

and that job it does marvellously. It's quick to start up and pleasant to operate – if you read a lot of news and blog sites which support RSS, this program will be very useful to you.

The project, like many in the Open Source world, is actively seeking contributions – current info on what is required is available on a regularly updated To-do list from the website.

SCUMM ENGINE

ScummVM

■ VERSION 0.5.1 ■ WEB www.scummvm.org

Back in the days when DOS was the king of PC operating systems, and the Amiga and Atari ST were still going strong, point and click adventure games were all the rage. LucasFilm games (later to be renamed as LucasArts) were behind many of the early classics; to save themselves the burden of rewriting each game from scratch, they created *SCUMM* (*Script Creation Utility for Maniac Mansion*, the first game), an engine which was specifically designed to interpret platform-independent code.

In this sense, *SCUMM* was like an early Java in that adventure games (the code of which described objects and actions etc) were compiled into a bytecode and then fed through the engine. This resulted in faster development and the ability to port games quickly to other platforms. After barely mentioning it in our Wine

feature in *LXF42*, Reader Dave Wickham suggested that we give *ScummVM* some coverage – it's an SDL-based *SCUMM* (and others) engine which runs on Linux, Windows, Mac OS and even PalmOS and Windows CE.

Compiling *ScummVM* from source should be relatively free of troubles providing that you have the SDL development libraries and headers installed. Note that the `--disable-mad` option doesn't appear to work correctly, so you may have to dig into the Makefiles and remove `-DUSE_MAD` in the DEFINES and `-lmad` from the LIBS options. Failing that, try the RPM supplied on our coverdisc.

ScummVM's interface is a simplistic graphical browser, in which game lists can be created and a few options tweaked – the command-line flags allow for a much greater amount of



Presumably he's waiting to mug a passer-by to get an present for Kayla.

settings to be modified. Right now, *ScummVM* will run *Monkey Island 1&2*, *Beneath A Steel Sky* (which was available along with *ScummVM* on last month's *Linux Format* coverdiscs) *Simon The Sorcerer 1&2*, *Zak McKracken*, *Sam & Max* and others with varying degrees of success. We had few problems with the titles we tested, and fortunately the engine offers save

states so the odd crash shouldn't be too painful.

A number of demos are available on the project's site if you want to check compatibility, and the coders and working hard on cleaning up the glitches. The remaining issues aside, it's an excellent way to enjoy the oldies without faffing around with the likes of *DOSemu* and co.

WEB SERVER

Cherokee

■ VERSION 0.4.6 ■ WEB www.0x50.org

In terms of UNIX Internet server software, the 'big three' are undoubtedly *Apache*, *Sendmail* and *BIND*. This trio has been around for many years – they're trusted, very reliable and have proven to be mostly tip-top on the security front. Because they're so widely used, though, they tend to be unwieldy and over-complicated to the learning admin and, equally, for a small server their impressive feature stats could maybe be considered to demonstrate overkill. We looked at *Monkey HTTPD*, a lightweight web server, back in *LXF41*; recently reader Dave Turner pointed us at *Cherokee* and asked us to give it a quick look too. If you've got your own requests for *Hot Picks*, please send them to the addresses at the beginning of this article.

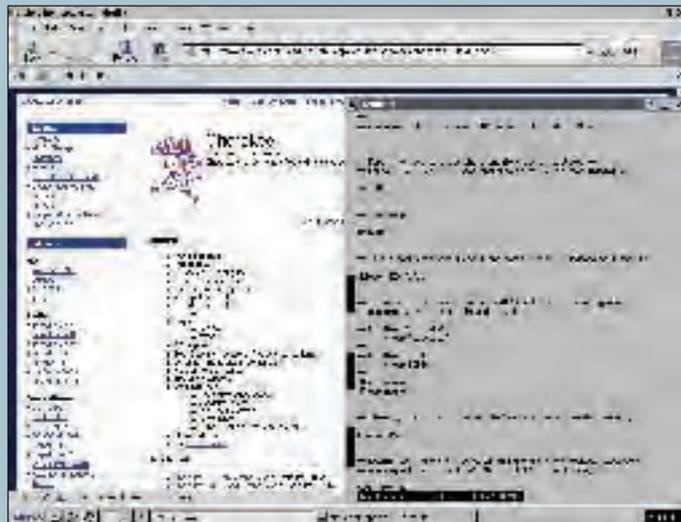
Cherokee aims to be a slimline and speedy web server which folds in many of the goodies that users have come to expect from *Apache*, while retaining a clean and accessible codebase. At only 348KB for the source, it certainly packs in a lot of bang for byte and should compile on virtually any system – there are few exotic dependencies, and the configure script automatically disables anything missing. Unless `--prefix` is changed, the binaries, libraries, development headers and config file are all dropped into the `/usr/local` tree.

Pleasingly, *Cherokee*'s config file resembles *Apache*'s very closely, with many of the same keywords and terse yet appropriate comments. (The developer warns users to make no assumptions in scripts etc. as new code and changes are added frequently.) In

its current state, *Cherokee* boasts CGI and partial PHP support, virtual servers, SSL, IPv6, NCSA/W3C logfile formatting, a plugin system and other touches which put it on par with *Apache* for most typical jobs.

So does its diminutive size help it meet the 'extremely fast' claims? Using *ab* for some benchmarking, we found *Cherokee* to serve static HTML

pages approximately four to five times faster than *Apache* on average – very impressive. Of course, this will vary a fair amount depending on configuration and other factors, but it's a good sign nonetheless. In all, if you're looking for a zippy, petite and straightforward web server with a decent range of features, *Cherokee* is a worthy choice.



Cherokee's site and config file – as exciting as things gets, visually.

INTERNET MONOPOLY CLIENT

Atlantik

■ VERSION 0.5.3 ■ WEB <http://unixcode.org/atlantik/>

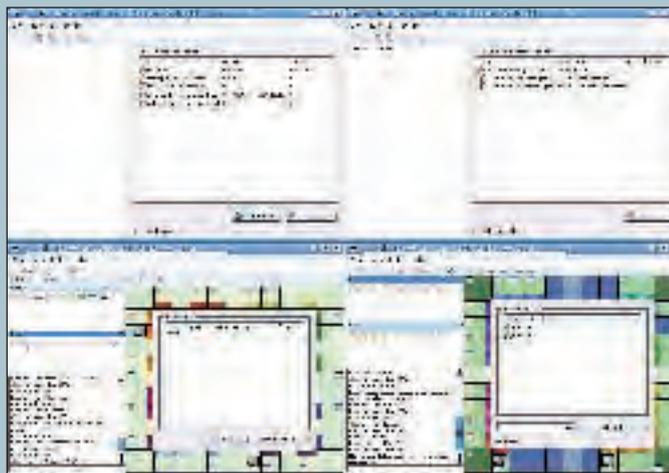
Question: what eats up six hours of your time, leads to endless squabbles, and from which no one person comes out feeling pretty good? Answer: a big family picnic. If that doesn't float your boat, *Monopoly* has the same effect. By far the best way to teach young children that making money is more important than anything else in the history of the Universe, Hasbro's *Monopoly* franchise continues on its successful quest to make people quip 'Where's Old Kent Road anyway?' and 'I bet today's RailTrack-owned stations are cheaper than that'.

Being a slow-paced board game (although admittedly top-notch fun when you've snagged Mayfair and have a hotel on it), *Monopoly* is ideally suited to online play. Consequently, various *Monopoly*

game servers have emerged, offering players from around the world a chance to thrash one another to a pulp. A client-side program is needed though, and *Atlantik* (formerly known as *KMonop*) is a KDE-base front-end and reader Craig James suggested that we cover it.

You'll need the relevant versions of *kdebase*, *kdelibs* and *kdegames* installed (along with their development libraries and header files) to build *Atlantik* from source. As of KDE 3.1 it has been rolled into the main games bundle, but development continues on the project's site.

Upon starting, you're presented with a server list of the *monopd* network which details where all the big games are ongoing. Servers are spread around the world, with players of varying ability battling away; as with *BZFlag* and the like, some servers are



A montage of shots: server and game selection, and middle of a match.

more responsive than others and it's often hard to find an open game, but if you've arranged with a friend then it's no problem.

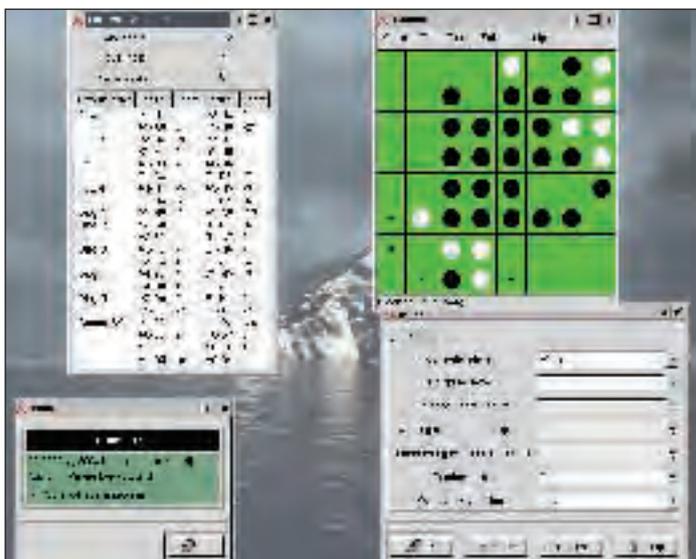
Atlantik offers variants of the main game, with the US and UK versions fully catered for, and players can seek out others with resources they need and offer to exchange money for property and vice versa. Otherwise, the play mechanics are as to be expected and *Atlantik's* vanilla visuals do the job decently enough.

Although *Monopoly* isn't everyone's idea of a fun night in, *Atlantik* is a first-rate network game client. If you enjoy boardgames of any sort, but don't fancy giving *Atlantik* a whirl, why not visit the website anyway? The author Rob Kaper is appealing for fans of other boardgames to put their ideas forward. His eventual aim is to organise a colossal Open Source server for all sorts of boardgames – a worthy effort indeed!

OTHELLO GAME

GRhino

■ VERSION 0.8.2 ■ WEB <http://rhino.sourceforge.net>



Ye Gods! It's the bright greenicity that's making me lose. And it makes GRhino a bit difficult to hide behind things when the boss appears...

Othello's slogan, from marketer

Mattel, pretty much sums up all human activities including post-row door slamming and jumping queues in supermarkets: "a minute to learn, a lifetime to master". The ostensibly simple black and white board game was developed by Mr Goro Hasegawa in 1971 as a variant of the traditional *Go*, and boasts major international tournaments. Even the Pope enjoys an odd match. Small *Othello* (and *Reversi*) desktop computer games have been available for some time, but for a real battle we're tempted by *GRhino*...

Described by the developers as 'targeted for experienced *Othello* players', *GRhino* aims to supply a tough AI engine and extra goodies for talented fans. You'll need relevant development packages for GTK 1.x and *gnome-libs* installed to build from source. A full GNOME 2 flavour is not yet available, but it should still compile on newer distros regardless.

GRhino's main game window includes no eye candy whatsoever; the marble texturing and silky-smooth rotation effects of *Iagno* aren't present,

but serious players won't be concerned with the lack of such frilly distractions. It's a shame there's no way to alter the default background colour though – staring at the garish bright green background when planning a move isn't pleasant.

But, to rectify, *GRhino* plays a mean game of *Othello* and the highest opposition level is a nightmare to compete against. Subtle alterations can be made to the AI settings through the Prefs box, while options are available to undo/redo the previous move and indicate visually the next possible moves. *GRhino* is also furnished with an openings database to demonstrate early play strategies – a sweet addition.

Even though it'll never win any awards for eye candy, *GRhino's* goal is to present the competent *Othello* player with a demanding challenge and that it does superbly (still, real-life networked play would be nice to see). Also, being small and easy to compile, it makes an ideal little desktop game to have around. If you become good enough, a night in with John Paul II could be on the cards. **LXF**

Kernel 2.6 is born!

The next evolutionary step in the development of the Linux kernel is now near release. **Richard Drummond** explains what it means for you and how to make use of it.

cover feature



Britons like to gamble on odd things, such as whether the temperature will top 100 degrees Fahrenheit or whether it will snow on Christmas day. While writing this article, I wondered if William Hill would accept bets on when the first stable release of version 2.6 of the Linux kernel will be made. You see, despite months of expectant speculation, that date is not yet known. Open Source development doesn't tie projects down to release schedules cast in stone: there's no marketing department to appease,

after all. Kernel 2.6 will be released when Linus Torvalds decides it's ready. The official line is that this will be some time in the fourth quarter of 2003, while, given Linus's tendency to aim for memorable dates, the smart money would be on Halloween, a date already hinted at. Not being a gambling man myself, and not currently being a UK resident, however, I didn't find out whether it was actually possible to make a flutter.

While we don't yet know when the 2.6 kernel will officially be released, we do have a fair idea of what it will look like when it gets here. This is because development on 2.6 is currently in feature-freeze. During this phase, no new features will be (or at least should be) added. The focus, >>

KERNEL 2.6



instead, is making sure the features that are already there work.

By the time this article is in print and on the news stand, it's possible that the first stable 2.6 kernel will already be here – but we doubt it! At the time of writing, the fourth test release of the 2.6 kernel has just been made, the fourth release aimed at stabilising the kernel and ironing out bugs, and there's still a lot of work to be done. For comparison, the previous stable release, kernel 2.4, took a full twelve test releases to stabilise it enough for it to be deemed fit for general consumption.

This testing phase has begun after seventy-five releases in the 2.5 development series since it was split off from the stable 2.4 kernel in November 2001. Since then the kernel source tree has grown by about 40 per cent in size as new features and new architectures have been added. With the caveat that some of what we say here may be subject to change as the

testing cycle continues, in this article we are going to examine some of these new features and discuss how to build and use a 2.6 kernel. Some of the changes in the new kernel require that you upgrade other parts of your system too, and some reconfiguration may be necessary, but, despite such issues, upgrading to 2.6 is something anybody with a modicum of Linux knowledge can attempt. It's definitely something that everybody can benefit from.

What's in it for me?

The last major update in the Linux kernel was the 2.4.0 release seems quite long ago – back in January 2001. The major theme of that version was scaling up. It offered numerous improvements for the enterprise user, including better support for SMP (Symmetric Multi-Processing) machines, allowing Linux to perform well on the kind of 2-, 4- and 8-way multiprocessor boxes typically used for large server loads. However, there was also a lot of good stuff for the desktop user in there too, such as the new DRM architecture for hardware-accelerated 3D graphics in X, the USB driver subsystem, and built-in ISA plug-and-play support.

Kernel 2.6 offers improvements and new features that should please all Linux users, not just in the server realm and on the desktop. With 2.6, one aim is to make Linux not just scale up better, but scale downwardly too. This includes widening the range of hardware that Linux will run on and offering better overall performance across the board, especially under heavy workloads. Several new

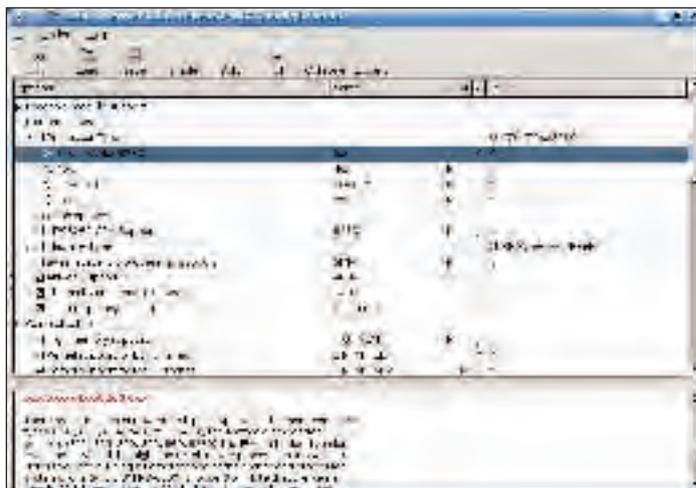
processor architectures have been merged in 2.6, including some embedded microcontrollers and some new 64-bit platforms (see the box All devices great and small). Support for SMP systems has been improved further and support for NUMA (Non-Uniform Memory Access) systems added. (The NUMA multiprocessing model attempts to address the memory bandwidth problems of SMP systems caused by all processors accessing a single pool of memory through a shared bus. In a NUMA system each processor has its own local memory, which it can access more quickly than non-local memory – hence 'non-unified!') Improvements for the embedded space begin not least with a greater degree of control of the kernel build process, particularly with regard to what can kernel features can be left out, enabling lower footprint kernels to be built when space is a premium.

Bigger, faster, more

The performance enhancements in 2.6 are many and varied, and should benefit all users no matter on what type of hardware comprises our systems and in what role we deploy Linux.

Firstly, the new O(1) task scheduler is much quicker at scheduling processes to run when a system is loaded with a large number of tasks (see the box *On schedulers*, above right). It also works more intelligently on SMP machines and attempts to be fairer when allocating processor time to tasks, trying to ensure that on heavily loaded systems that no process is starved for a bite at the processor.

Kernel 2.6 offers some improved GUI configuration front-ends to replace the unnecessarily ugly TK-based GUI in earlier kernels.



All devices great and small

The many processor architectures supported by Linux

Linux was originally designed for use on standard x86 PCs, without any clear intention that it would ever work successfully on any other architecture. Portability was never considered, and Linus Torvalds even stated early on that it would never run on anything else. Lucky for us and the computing world that history has proved him wrong!

The first port to appear was of kernel 0.99pl9 to the Amiga in 1993, and, not long after, the Atari ST (support was later added for other Motorola 680x0-based hardware such as pre-PowerPC Macs from Apple and early Sun boxes). This work was done outside of the main kernel

tree, however, and wasn't merged in a stable release until kernel 2.0. The official kernel began to support non-x86 platforms with the release of 1.2.0, which could be built to run on Alpha, MIPS and SPARC processors as well as the ubiquitous x86. Since then, Linux has picked up more architectures with every major release.

Supported architectures

Kernel 2.6 brings the number of supported architectures to 20. New conquests are the 64-bit PowerPC architecture (such as IBM's iSeries and

pSeries servers and – eventually, fingers crossed – Apple's new G5 Mac) and AMD's x86-64, their 64-bit replacement for the standard x86 (Intel) architecture. Other ports have been merged too. 2.6 incorporates the uClinux project (www.uclinux.org/), which supports several MMU-less processors typically used in embedded devices (such as the 68K-derived Dragonball and ColdFire series from Motorola). Usermode Linux is also now part of the 2.6 tree. This isn't a physical architecture, but it lets you compile Linux as a user-space executable which can be run to create virtual Linux environments – great for debugging

purposes or running services in a secure environment. Other changes include the MIPS-64 port being folded in with the main MIPS port and similarly the 64-bit and 31-bit (S/390) zSeries ports have been combined. Meanwhile, support for older ARM processors (the 26-bit chips such as those found in the Archimedes and other RiscOS machines) has been split off from the main 32-bit ARM port.

The sheer breadth of hardware supported by 2.6 is staggering. From the same source tree, you can build kernels that will run on anything from a Palm Pilot to the latest zSeries mainframes from IBM.

Secondly, support for threading has been improved, for instance, allowing threads to be created and destroyed much more quickly. Much of the changes have been wrought to support Ulrich Drepper's new NPTL (Native POSIX Threading Library), which implements a more robust and efficient userspace threading environment, supporting POSIX signals and thread-local storage. This is something that enterprise users have been clamouring for for some time, and it's easy to see why. Tests conducted by Drepper showed that with the NPTL and the kernel changes it is possible to create and destroy 100,000 threads in *two seconds* on a standard x86 box stuffed with 1GB of RAM. On an earlier kernel without the threading modifications, this would have taken about *15 minutes*. The NPTL will eventually be merged with GNU's *libc*, but it is currently available from <ftp://people.redhat.com/drepper/nptl/>.

Thirdly, Linux's block I/O layer has been re-written for improved efficiency, and users can now choose from several block I/O schedulers the one that is most suitable for their target system and workload. The block I/O subsystem is the kernel layer that interfaces between the device driver for a block device such as a disk and the filesystem handler that provides high-level access to the data stored on that device. Not surprisingly it is a critical component in determining disk throughput. Several tasks can make request to access a disk at the same time, and it is the job of the I/O scheduler (also known as the I/O elevator) to determine the order in which these requests get fulfilled.

The new default scheduler is the anticipatory I/O scheduler. Its name derives from the fact that when a process issues a request to read from a block device, this scheduler performs that read and then waits a few milliseconds – even if there are requests from other tasks queued – anticipating another read request from that task to follow. By waiting, instead of fulfilling a different request which would probably access a different physical location on disk, the anticipatory scheduler reduces the number of unnecessary disk seeks performed and thus increases overall disk bandwidth.

On schedulers

The new O(1) scheduler in detail

Linux 2.6 features the new O(1) scheduler, largely the work of Ingo Molnar. If you've not studied computing science before, then the term 'O(1)' probably won't mean much to you. This is a notation used for classifying the performance of an algorithm in relation to the size of the input being processed. The 'O' here stands for order. If an algorithm is classified as O(n), for example, this means the time taken for it to complete grows in direct proportion to the size of the input, n; if the input doubles in size, the algorithm will take twice as long. If an algorithm is O(n²), then the time taken varies

with the square of the input size; if the input doubles, the time taken quadruples. The holy grail is to find algorithms which complete in a constant time, no matter how big the input size. This is known as O(1).

Thus with the O(1) scheduler, on uniprocessor systems the work required to schedule a new task to run is roughly the same no matter how many tasks are running. For systems lightly loaded with tasks, the new scheduler performs comparably to the scheduler in 2.4, but for systems running large numbers of processes, scheduling is much quicker than in 2.4.

The new scheduler is also much better at scheduling interactive tasks. It does this by collecting statistics on scheduled tasks to try and guess which are interactive and which are batch tasks. Simplifying things slightly, it assumes tasks which never voluntarily give up the processor and sleep are batch tasks that will use all the CPU time given them, and tasks which do sleep occasionally (such as when waiting for user input) are interactive tasks. Based on these estimates, it gives a priority bonus to tasks based on their level of 'interactivity' so that interactive tasks can be scheduled to run and perform well even under heavy system load.

Two other schedulers are also currently included, the deadline I/O scheduler and the no-op I/O scheduler. The former tries to ensure that all I/O requests are fulfilled within a specified deadline, and generally offers performance comparable to the anticipatory scheduler when only a single process is accessing the disk at any one time. It is a less complex and hence smaller, so may be preferable in certain environments. The no-op scheduler, on the other hand, is a very basic scheduler, which is suitable for memory-based block devices rather than disks. Other I/O schedulers are being developed, implementing different scheduling policies, and may be merged before 2.6 is officially released.

Low latencies

The performance enhancements discussed so far should benefit all users, but their affect will be greatest felt on systems that are worked hard, which generally means servers rather than desktop machines. However, there's good news in 2.6 for desktop users, too. Historically, Linux has not been good at jobs, such as playing multimedia streams, which require tasks to be scheduled with minimum latency in response to events. This has been due to the monolithic design of the kernel itself. It was possible for tasks calling on kernel services to get stuck in kernel space for relatively long periods of time, thus not allowing any other task to be scheduled. This problem has been attacked in two ways in 2.6, basically merging work



that has existed outside the main kernel tree for some time.

The first is Andrew Morton's series of patches that reduce scheduling latency. The second is Robert's Love kernel pre-emption patch. The latter permits the optional building of a kernel that is pre-emptible; that is, a task executing in kernel space can be pre-empted or put to sleep and a higher priority task scheduled to run. These two changes mean that tasks which need to be scheduled regularly with low latencies, such as a task that regularly refills a soundcard's buffer with a waveform to be played, are less

The new kernel will bring support for 64-bit architectures, eventually including even Apple's new G5 Mac.



KERNEL 2.6



likely to be delayed waiting for some other task to finish executing in kernel space. Linux is still not a real-time system – no guaranteed response times are provided – but for the desktop and in many embedded applications this is a big step forward.

Devices and power

Kernel 2.4 featured a much-improved device architecture over previous kernels, with support for resource management and an integrated infrastructure for PCI, PC Card and ISA plug-and-play devices. 2.6 takes this much further, with a new unified,

any PCI devices as children of those, and so on. A logical 'system' branch in this tree is created for system devices like the CPU, interrupt controller and timer. A 'legacy' branch is also created for devices – that are on legacy busses and which can't be probed – such as non-plug-and-play ISA devices. This tree is used by the power management system to ensure that devices are powered up and powered down in the correct order. It is also exported to userspace with a new pseudo filesystem, called *sysfs* (see the box *SYSFS*, far right).

The new device model lays the foundation for full power management of devices. On x86 platforms, for example, two power management schemes are now fully supported, APM and ACPI (Advanced Configuration and Power Interface). The former uses BIOS services to take care of power management and handle putting the system to sleep.

ACPI, on the other hand, puts power management fully under the operating system's control and is featured on all modern x86 systems. ACPI is used to discover and monitor power-related devices in a system, such as battery and mains power state, power buttons, and even temperature and fan status. ACPI is also necessary to enable Hyper-Threading on recent Pentium 4 chips. Thanks to updated ACPI support in 2.6, and the interface it exports via *sysfs*, it is possible to switch systems into lower power states or put them to sleep entirely. Desktop systems typically don't support this, but they do often have a CPU 'throttling' control, again accessible via *sysfs*, which can be used to slow down the processor – handy for letting an overheating system cool down. Plenty of userspace tools are available which make use of ACPI for things like monitoring battery level. The ACPI daemon (see <http://acpid.sf.net>) can be used to execute commands or scripts when ACPI events occur, such as the power button being pressed or a change in mains power state.

Also related to power management, kernel 2.6 features CPU frequency scaling on supported processors and systems, such as x86 laptops powered by 'mobile' processors and certain iBooks and PowerBooks. Lowering a CPU's operating frequency, consumes

less power and thus extends battery life on laptops. This feature is again accessed – you guessed it – via *sysfs*, and client software can be installed to automatically put it to use. For example, the *cpudyn* daemon (see <http://mnm.uib.es/~gallir/cpudyn/>), will switch your CPU to a low frequency state when your system has been idle for a specified time and can also automatically spin down your harddrive.

Better driving

We've looked at some of the main architectural changes in kernel 2.6, but the update provides many other advancements as well. Several driver subsystems have been substantially reworked, driver support has been widened, and many driver projects that have previously existed outside the main kernel have been merged.

The IDE subsystem was a cause of much controversy early on in the development of the 2.5 series. The initial re-implementation of the IDE code was dropped and it looked like the much needed overhaul wouldn't be seen in 2.6. Alan Cox took over and cleaned up the IDE system, however, and with recent work from Bartłomiej Zolnierkiewicz enabling taskfile-based I/O, IDE support in 2.6 is in pretty good shape. Tagged Command Queueing (TCQ) is now supported on certain drive and controller combinations and improves disk throughput by allowing a controller to issue a request to a disk before the previous request has completed. Support for some Serial ATA controllers including chipsets from Promise, Intel and VIA has been added, and it is now possible to use ATAPI CD writers without resorting to the SCSI emulation layer.

Multimedia support in 2.6 gets a boost, as several important external projects have been merged. Firstly, the ALSA project, the next-generation audio architecture for Linux, has been brought into the official kernel tree. This provides a much cleaner and more modular system for driving soundcards than the existing OSS (Open Sound System) drivers. It also introduces wider hardware support and powerful features such as a plug-in system and a flexible in-kernel sequencer. We covered ALSA in detail in *LXF44*, so see that issue and the

“Be aware that while the 2.6 kernel is in the test phase it is still beta-quality. It may work, but it may also cause you problems.”

object-oriented device model. All devices, no matter their bus type, are based on a common device object, which defines services and attributes applicable to all device types, such as device probing, power management and reference counting. The different subsystems add to this with methods and attributes appropriate to their behaviour. Not only is this new model conceptually cleaner and more elegant, it allows proper power management of devices for the first time.

As the kernel boots and discovers devices attached to the system, it links all these devices into a single logical tree that reflects the relationship between these devices. At the root of this tree will be your system's north-bridge, with any PCI buses as children,

No limits

Scaling up to some really large numbers in 2.6

Thanks to support for 64-bit access to block devices, Linux can now make use of very large disks. The limit has upped from 2TB to 16TB on 32-bit architectures (1TB or TeraByte is 1024GB, roughly a thousand billion bytes) and to 8EB on 64-bit architectures (1EB or ExaByte is roughly a million TeraBytes or 10¹⁸ bytes – that's a 1 followed by 18 zeroes). The latter limit that shouldn't cause any problems for the next decade or two!

Kernel 2.6 adds support for Intel's Physical Address Extension (PAE) which ups the amount of physical memory

addressable on the x86 architecture from 4GB to 64GB.

User and group IDs are now represented by 32-bit numbers rather than 16-bit numbers. Thus kernel 2.6 system allows more than 4 billions user accounts and 4 billion groups. The previous limit was 65,336 of each.

Process IDs are now allocated using 30 bits rather than 15 bits meaning that processes can be numbered up to roughly 1 billion before PID numbering wraps around and begins at 1 again. The previous limit was 32,767.

ALSA website (www.alsa-project.org) for more information. Also merged in the 2.6 kernel is the LinuxTV project (www.linuxtv.org) which includes drivers for DVB devices, such as digital TV receiver cards. Finally, the Video4Linux system, which handles conventional video capture devices, has been updated to support the V4L2 API – thus providing, for example, an improved interface to devices with their own on-board hardware compression.

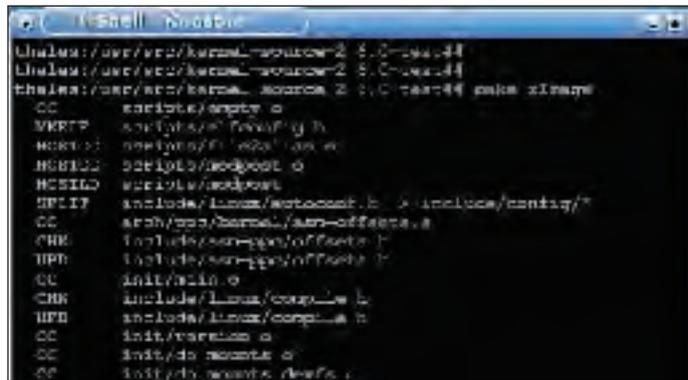
The kernel framebuffer has also had a much needed overhaul in 2.6. While functionally similar to the existing implementation, the framebuffer system has now been completely detached from the kernel console layer allowing the framebuffer to be used to drive a graphics card without forcing you to put a virtual terminal on that display too. Many of the framebuffer drivers have been cleaned up too, and feature improved hardware acceleration and monitor detection. Currently the framebuffer still needs some tidying up before 2.6 is officially released.

The `lm_sensors` project (at <http://secure.netroedge.com/~lm78/>) has been partially merged with 2.6, augmenting the kernel's existing support for I2C devices with drivers for a range of system health and monitoring devices, such as the fan speed and temperature sensors commonly integrated with motherboard chipsets. Related is inclusion of support for IPMI, Intel's Intelligent Platform Management Interface (www.intel.com/design/servers/ipmi/), a system which allows remote monitoring of the health of servers with compatible sensors.

We've don't have the space to cover all the improvement and enhancements in 2.6. For further details, see take a look at Joseph Pranevich's *Wonderful World of Linux 2.6* www.kniggit.net/wwol26.html or, for a more technical view, Dave Jones's Post-Halloween Docs at www.codemonkey.org.uk

Trying it yourself

So you've heard about all this great new stuff in the 2.6 kernel, and, at this point, you're probably thinking that you really want to give it a go yourself. Good. That's what we're going to look at next. Before you jump in with both feet however, stop and think for a



```

Linux: /usr/src/kernel-source-2.6.0-rc4-44
Linux: /usr/src/kernel-source-2.6.0-rc4-44
this: /usr/src/kernel-source-2.6.0-rc4-44 make xImage
CC      scripts/empty.o
MKDIR   scripts/objtool
MKDIR   scripts/objtool
MKDIR   scripts/objtool
MKDIR   scripts/objtool
MKDIR   include/linux/asm-generic
MKDIR   arch/arm/kernel/asm-offsets
MKDIR   include/linux/asm-offsets
MKDIR   include/linux/asm-offsets
CC      init/init.o
CC      include/linux/compile.h
MKDIR   include/linux/compile.h
CC      init/ramdisk.o
CC      init/do_mounts.o
  
```

The kernel build system has been overhauled in 2.6 and is now much quicker and cleaner.

minute. Building the kernel from source yourself isn't really that difficult an exercise, but you do need to be careful.

If you've never built a kernel before, you should at least read the Kernel HOWTO from the Linux Documentation Project (in the 'Help' section of this month's LXF coverdiscs or www.tldp.org). This guide hasn't yet been updated to cover the 2.6 kernel, but the majority of the information it provides is still relevant, and it provides an excellent grounding in the concepts you need to grasp to successfully build a kernel.

Next, you should be aware that while the 2.6 kernel is in the test phase it is still beta-quality software. It may work for you, but it may cause problems. In any case, it's certainly not ready for running on production systems. Even when the first full 2.6 release is made, reliability will not be as high as with the current stable 2.4 series simply because it hasn't had the widespread testing and volume of bug reports that 2.4 has been through over the last two-and-a-half years.

Finally, upgrading to 2.6 will cause some compatibility problems (we'll look at a few later). You may need to upgrade many of the user-space tools that accompany the Linux kernel to be able to use it on your system. One particular difference in 2.6 is the change in the kernel module loading system. A new set of tools is necessary, called *module-init-tools*, and configuration of *modprobe* is subtly different. More on this in a moment.

The latest kernel source code can always be obtained when you need it from www.kernel.org or one of its many mirrors. Follow the link on the front page of this site for the latest 2.6 release; or, alternatively you can go

straight to www.kernel.org/pub/linux/kernel/v2.6/. Grab the latest kernel 2.6 source tarball.

Now unpack the kernel. Some people use the `/usr/src` directory to store kernel trees, but there's no reason you can't use your home directory. If you opt for the latter, you don't need root permission to build the kernel (although you still need to be root to install it).

After unpacking the kernel, before you do anything else, browse through the files `README` and `Documentation/Changes`. The `Changes` file tells you what tools you need to build the kernel and what user-space tools are required when running the kernel. Ensure that all your tools are up to date. For most architectures version 2.95.3 or better of the C compiler GCC should be used. A GCC release from the 3.x series may also be used, but newer GCC releases haven't been so well tested for kernel building, so you may discover problems. Having said that, GCC 3.3.1 works fine for me on x86 and PPC.

Module-handling tools

The chances are, if you have a reasonably modern distro, you will have the necessary tools already. One exception will be the new module handling tools, *module-init-tools*. Get the latest release (currently 0.9.13) from www.kernel.org/pub/linux/kernel/people/rusty/modules/ and install it before building your kernel. The following series of commands will build and install *module-init-tools* from source.

```
./configure --prefix=/
```

```
make moveold
```

```
make && make install
```

Note that the new tools offer

SYSFS

Accessing driver internal from userspace

Kernel 2.6 features a new pseudo filesystem called *sysfs* which is used by device drivers to expose their innards to userspace. As well as providing information on devices attached to your system and their associated drivers, many drivers can be tuned or allow their behaviour to be modified by their *sysfs* interface. In many ways this a lot like the existing `/proc` filesystem, but *sysfs* provides a simpler interface. Each file in *sysfs* exports a single attribute and is guaranteed to use only ASCII characters. Thus you only need simple console tools to read and write values to *sysfs*.

The standard mount-point for *sysfs* is the directory `/sys` at the root of your filesystem. Create this and mount *sysfs* manually with `mount none /sys -t sysfs`

This should also be done somewhere in your startup scripts. Track down the script that mounts `/proc` on your system and add the above line immediately afterwards (it's `/etc/init.d/checkroot.sh` on Debian systems and `/etc/rc.sysinit` on Red Hat). You should also add the *sysfs* mount-point to your `/etc/fstab` file as follows:

```
none /sys sysfs defaults 0 0
```



KERNEL 2.6



backwards compatibility to the old *modutils* tools: they will fall back on using the old tools when running under a 2.4 kernel. The line **make moveold** on the previous page renames your old **insmod**, **modprobe**, *etc* commands to **insmod.old**, **modprobe.old**, *et.* so that they are available as a fallback to the new commands. You only need to do this the first time you install, however. Also note the new module loading system requires a subtly different configuration (see box, *Loadable modules*, below).

Debian users who are tracking *sid* (the unstable Debian tree) can save themselves some effort by doing **apt-get install module-init-tools**. This will also generate a *modprobe* configuration suitable for your setup.

Construction zone

Building a kernel is conceptually quite simple. The first step is to configure the build process – that is, decide what features you need, what is to be built-in, what is to be a loadable module, and so on. Once that's done, you then build the kernel image itself, the loadable modules and install everything. The kernel build system

Loadable modules

Configuring module loading in 2.6

The way the kernel handles loadable modules has been overhauled in 2.6. Principally this was to reduce the chance of race conditions during module unloading (to avoid the possibility that a module could be used when in the process of being unloaded), but an added bonus is improved hotplug support. However, you do now have to explicitly enable module unloading when building a kernel if you want this feature – and most people will. Another change is that modules now have the filename extension *.ko* rather than just *.o* to underline the fact that they are kernel modules.

Kernel 2.6 requires a new set of userspace tools to handle module loading, called *module-init-tools* (see the main article). Although **modprobe** can be used without any configuration, if you want to supply default options for modules or use module aliasing to enable the demand-loading of modules by the in-kernel module loader, *kerneld*, the you do need to configure it. Previous **modprobe** was configured via the file */etc/modutils.conf*. This has now changed to */etc/modprobe.conf*,

and fewer configuration options are now supported than before (see *man modprobe.conf*). Notably, you can no longer have aliases to other aliases. Luckily, a script is provided in the *module-init-tools* package to ease the transition. This will create an appropriate *modprobe.conf* file from your old */etc/modules.conf*, and can used as follows:

```
./generate-modprobe.conf
/etc/modprobe.conf
```

If you use the device filesystem (*devfs*) to create */dev* nodes on the fly, then you will also want to copy the *modprobe.devfs* file supplied with *module-init-tools* to */etc*. If module auto-loading is enabled in your *devfs* configuration (as it will be for most), then, when a process tries to access a non-existent */dev* node, *devfsd* calls *modprobe* with the name of that node. The *modprobe.devfs* file contains a list of aliases which will map between common */dev* nodes and the modules that need to be loaded to create those nodes. Thus you need this file for demand-loading of modules to work fully with *devfs* under 2.6.



Performance improvements such as the new O(1) scheduler and kernel pre-emption, will allow smoother handling of multimedia streams in 2.6.

has been significantly overhauled in 2.6, thanks to work done by Roman Zippel. Not only is it quicker and easier to use than before, there are several new GUI-based configuration front-ends to replace the old and ugly *Tk*-based one.

With the new system, the series of commands in a typical kernel building session would now look something like:

```
make menuconfig
make all
make install && make
modules_install
```

First note that there's no longer a **make dep** step. This was previously required for the build scripts to figure out dependencies between kernel components, but now you just need to issue a command to build your desired image (here we used the default **make all**) and everything will be done automatically.

Let's look at this process in more detail. The real purpose of the configuration stage is to generate a file called *.config* which is used by the build system to decide what features to include in the kernel that it creates. To do this execute the appropriate *make* command to launch your preferred configuration front-end and select the options you need. Above, we chose **make menuconfig**, which uses a tree of console-based dialogs

to do the job. Some of the other configuration targets are:

make config	Standard interactive text-based interface
make xconfig	New Qt-based interface (requires Qt libraries).
make gconfig	New GTK+2.0 interface (requires GTK+ and Glade).
make defconfig	Use default config for your architecture (non-interactive).

Note that you may first supply an existing *.config* file from a previously compiled kernel to cut down on the work you need to do – by simply copying it into the root folder of the kernel source tree. You can then issue **make oldconfig** to pick up and configure any new kernel options. This isn't recommended under 2.6 if you are using an old 2.4 *.config* file for this because of subtle changes in many of the configuration options in 2.6. Use one of the full config front-ends instead.

When configuring your kernel, make sure that you include at least the minimum amount of services required to boot your system as built-in modules. In particular, remember drivers for your disk interface and disk, your root filesystem, and

whatever drivers are required to support your input devices, such as keyboard, mouse and, for systems without a VGA console, a framebuffer driver (you also now need to explicitly enable the framebuffer console as well as a chipset driver). Include TCP/IP networking as a built-in, but you only need a network-card driver built-in if you boot from a network. Note that for the dialog-based config systems, the layout of the various configuration sections has been substantially re-ordered and should make navigation simpler.

In the example on the previous page, we used the default **make all** target to build the kernel image and modules. This generates the default image type for your architecture (which is **bzImage** for x86 and **zImage** for PowerPC). Other build possible build targets include the following:

- make vmlinuz** Build the plain, uncompressed kernel image (needed for, eg Mac bootloaders).
- make zImage** Build a compressed kernel image.
- make bzImage** Build a 'big' compressed kernel image.
- make modules** Build the loadable modules.
- make rpm** Build the kernel and modules and package them as an RPM.

To install your new kernel, execute

```
make install
and to install the loadable modules, do
make modules install
```

The former should work for most, especially if you have a suitable **installkernel** command present on your system (Debian and Red Hat users will). Alternatively, you can install the kernel image manually (though you will still need to use

```
make modules install
On x86 systems, simply copy the compressed kernel image (found in arch/i386/boot/) and the System.map file to your system's boot folder. For example, as root enter
```

```
cp arch/i386/boot/bzImage /boot/vmlinuz-2.6.0-test4
cp System.map /boot/System.map-2.6.0-test4
```

Now add your new kernel to your

Safe as houses

New security features

Kernel 2.6 now includes the kernel cryptographic API, which was merged to support the IPsec Virtual Private Networking system that has also been integrated. The various cryptographic modules supplied with the kernel and which implement a range of common algorithms can also be used with the cryptoloop device, now also a kernel feature. This can be used for mounting encrypted filesystems, and is a great security feature for laptops.

Also new is support for different security models besides the familiar Linux security model, and the NSA's Security-Enhanced Linux project has now been merged. SE Linux (see www.nsa.gov/selinux/) introduces mandatory access controls to Linux which boosts security by restricting users and tasks to a minimum set of necessary privileges.



Sysadmins can also fine-tune security policies by taking advantage of the **POSIX Access Control Lists (ACLs)** now supported by the main Linux filesystems.

ACLs allow the introduction of much more finely grained access permissions on files over the standard owner/group/world model employed in Linux.

bootloader configuration and reboot. With GRUB, for instance, you should add an entry similar to the following to your `/boot/grub/menu.lst` file. Remember to change the **root** command and the **root=** kernel option to specify your own root filesystem.

```
title Kernel 2.6.0-test4
```

```
root (hd0,2)
```

```
kernel /boot/vmlinuz-2.6.0-test4 ro
root=/dev/hda3
```

Possible problems

If all went well, then your new system should come up smiling as per usual – albeit now powered by a 2.6 kernel. However, as we said, some of the changes in 2.6 may require that your system be reconfigured, so problems may occur.

One potentially troublesome area is with loadable modules. The new module tools and configuration format may cause incompatibilities, and the names of some modules have even changed. In particular, the lack of support for aliases to other aliases in `modprobe.conf` will often cause issues with some systems. For example, if you were already using ALSA on your system, it would be quite common to configure the loading of ALSA modules in `/etc/modules.conf` with

```
alias snd-card-0 snd-cmipci
```

```
alias sound-slot-0 snd-card-0
```

This won't work in 2.6 and will need to be substituted in `modprobe.conf` with

```
alias snd-card-0 snd-cmipci
```

```
alias sound-slot-0 snd-cmipci
```

Also note the AGP driver in 2.6 has been split up with separate chipset drivers and this means that XFree86 will no longer be able to enable AGP support for your graphics card without some help. Find the name of the AGP driver for your chipset and add the following to your `/etc/modprobe.conf`, substituting the name of your driver instead of `via-agp` (which is for VIA chipsets).

```
install agpgart /sbin/modprobe --
ignore-install via-agp
```

Red Hat systems may have more serious problems, and hot-plugging of modules may not work at all. If this happens, edit the file `/etc/rc.sysinit` and replace all instances of `/proc/ksyms` with `/proc/kallsyms`.

Transitioning a system to the 2.6 kernel from 2.4 is not going to be without teething troubles. As potential problems are identified, solutions will be made available on all the usual support sites on the Internet – that's the beauty of Open Source software. Hopefully you will find the tweaking necessary to run 2.6 smoothly is justified by the new features and performance that it offers. If you are wary of upgrading, though, it won't be long before distros are updated with 2.6 kernels. Then you can make use of the advantages that it offers without any heartache. [LXF](#)

DEVELOPER PROFILE ZoneMinder

Don't overlook physical security! Nick Veitch chats to ZoneMinder developer Philip Coombes.

ZONE MINDER FACTS & FEATURES

- Uses PHP, Perl and C++
 - Makes use of SQL database
 - Supports V4L capture devices
 - Intelligent recording
 - Plenty of customisation potential
- www.zoneminder.com

You can even use the system to monitor garden wildlife – focus on a birdfeeder or watch hedgehogs nocturnally frolicking on the lawn.



Do nasty things lurk around your house at night? Do your fish need watching closely? Would you like to be able to look around your house when you're not there? The excellent *ZoneMinder* can help you.

Conceived by Philip Coombes after his garage and those of neighbours were burgled, *ZoneMinder* is actually a suite of tools and applications to provide intelligent CCTV monitoring – you just need to supply a camera and a box to attach it to.

Linux Format got to speak to author, Philip Coombes, about *ZoneMinder* and open source development in general.

LXF: The first question is: Why?

PHILIP COOMBES: I wrote it because my garage got broken into. Some people did the whole street. They took some nice wine and some power tools, but not much else really. It occurred to me that if I'd had a couple of cameras it would have been easy to see who it was.

The cameras were easy to come by but the software wasn't, so I thought I'd have a go. It mixes a lot of different disciplines – real time programming, web programming and database programming – so I thought it would be a good exercise as well.

LXF: A classic 'scratch your own itch' tale. Why did you decide make it open source?

PC: There are two reasons really. One is that I use open source programs, so it's nice to adopt that sort of philosophy, and I'm not

ZoneMinder supports all video for Linux devices, so pretty much any video input that you can get working with Linux, from USB webcams to ordinary video cameras connected to capture-cards.

The cunning part of the software, and from where the name derives, is that it can intelligently capture only the action that you require. The incoming image is divided into user specified zones. Some of these can be designated as of no interest (a tree that might wave about in the wind for example) and others can be giving critical importance (a doorway perhaps). The *ZoneMinder* daemon runs in the background comparing

incoming images. When a notifiable event occurs, the video is captured to disk so you can review the incidence in question at your leisure.

There are numerous C++ utilities and Perl scripts to facilitate the setup of the system, but most of the user interaction is through a web-based interface written in PHP. If you install the optional recommended Java streaming client, you can even view camera images in real time.

While configuring the software may take a little time and patience, once it's set up, it more or less takes care of itself. More interestingly, it has already identified some thieves, and is helping police with their enquiries.

just trying to catch the people who broke into my garage, I want to help other people.

Also, video is one of those areas that is very hardware dependent. There are so many different configurations and types of devices. You could use USB cameras, ordinary cameras into a capture card – I don't have the resources or time to test all this equipment so this is a good way of making sure it works – I have everyone else testing it for me.

LXF: We noticed that the software does a lot more than you originally needed it to. Was that something that you sat down with the idea of doing, are there bits and pieces that have been added over time?

PC: Both really. I have added quite a lot on since I started it. It was a lot simpler initially. People suggest to me things that would be useful for them, and if I think these things would be useful in general then I put it in. When I started the project, I did start with the idea that it would be easy to extend, so I didn't confine it to a rigid structure to start with.

LXF: Is that something learnt from years programming experience?

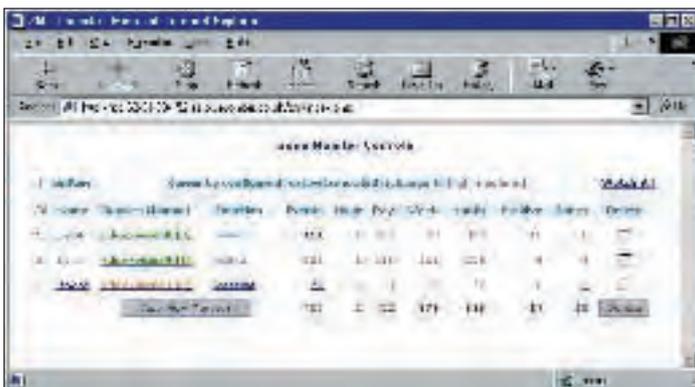
PC: Yes, pretty much. I find it's always easier to invest a bit of time to start with and cut down on the rework later on, rather than write something that just does the job, but you later have to re-architect to add something else. I know that's the opposite of the current trend to 'extreme programming', but that's my habit.

LXF: How quickly did you get feedback from the initial release?

PC: Fairly quickly. It started off as a bit of a trickle, but I did get some good initial feedback. Some of the people who gave me comments at the time are still helping now. Some people sort of shadow you, which can be useful.

LXF: How would you say the ratio between unwanted noise and useful feedback worked out?

PC: I think it's eighty or ninety percent useful feedback. Most of the noise is from people who haven't read the documentation. I've not had anything abusive or anyone just wasting my time.



LXF: You mentioned *ZoneMinder* had been used in evidence in a case. I guess that was the whole idea, but did it surprise you?

PC: I suppose it did in a way; although the guy who did it had got in touch with me beforehand, and that was always his intent. That was the plan, and the plan did work. He has said I can put the evidence on the website. I didn't want to jeopardise anything, so I haven't put it up yet. He has sent me MPEGs of the 'act'.

LXF: Was it good?

PC: Well, he showed the police and they recognised them straight away, and took them into custody almost immediately. I imagine in a lot of cases you wouldn't get that positive a response.

LXF: The police can often draft a list of suspects based on circumstances, because a large percentage of crimes are committed by people already known to the police.

Whether the video evidence is admissible or not in court, if they have an idea who did it, it helps!

PC: People use it for different purposes too. Watching other things. One camera I have pointing at a fishtank, one looks out for foxes. You can get underwater cameras, I was thinking of putting one in my pond!

There is another guy who is going on an expedition to Sumatra or Java or somewhere. It's a proper scientific expedition, but they don't have a huge budget so they can't afford pro-quality gear. I think the idea was he was going to use an Infra red camera attached to a laptop to track and locate rare tigers. I would never have thought of that.

LXF: That's often the way. You may start out with some simple goal,

but other people will pick up on it and use it in some other way.

PC: Someone got in touch with me today. They want to use it in some security app and wanted to know if I minded them forking it. It would still be under GPL, but there were some things they wanted to do. If there's stuff like that fed back in again, and get more of a community working, then so much the better. That's something I'm going to have to spend a little time on, getting CVS servers and admin stuff set up.

There are lots of different threads – there's developing it, promoting it, and then I suppose the admin side of things – from something that you do for an hour or so a couple of nights a week it soon starts to consume a lot of time.

LXF: From a structural point of view, the mix of languages is an interesting solution. PHP, Perl, C++. Why did you settle on that mix?

PC: Well, my experience is C++ going back a few years. Writing something that needed to be high performance for the video gathering and analysis, it was natural for that. Having said that, writing any sort of communications stuff, dealing with protocols is such a pain. Perl was so much easier – you just plug something in and off it goes. And it interfaces well with the database.

PHP, well if I'd started using it a year earlier I may have done away with Perl as well. I couldn't bring myself to do the web stuff in Perl, PHP was so much easier, and it was a learning thing. There are a few problems with the Java app used for streaming as well, so that's on my list, to write one of those.

LXF: What was the hardest bit of the code to write?

PC: The easiest bit was the C++. It

was obvious to me. The hardest was probably the configuration to some extent. Getting the configuration to work across all those different parts, although that has been superseded in later versions by including most of it in the database. I think it's probably that – the maintenance of a common core across three different technologies.

LXF: Any changes on the horizon?

PC: One is to do a PVR-style recording facility, so it just records continuously and cycles out old footage. Although that isn't something it was designed for, it's something a lot of people have asked for and it's fairly easy to do.

My roadmap also says that in two releases time I have to make it Multilanguage. I'm not going to do the translations though! But it will be a challenge to create the framework that will allow that to happen. The PHP on the web is one thing, the stuff emitted from the C++ applications is something else, and then there's the Perl scripts as well. Creating a framework for that might be tricky. The configuration framework for all of that was hard enough. The other thing is enhancing the wireless capabilities for phones. With the language, those are my next three things.

The phones seem simple but I only have two or three devices to try it on myself, so I'm planning to use some Open Source browser detection capability to help with that.

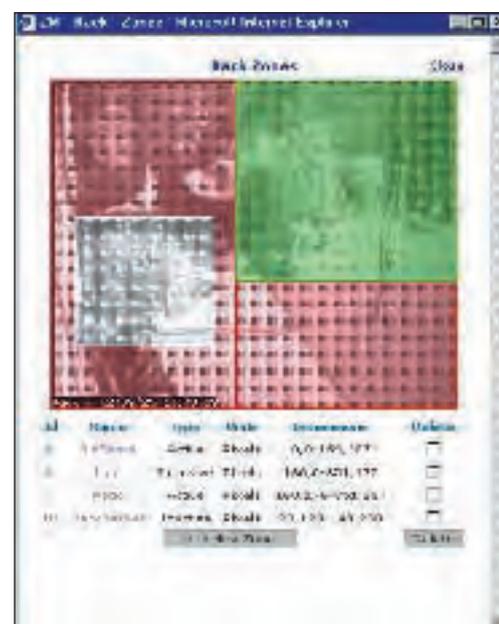
LXF: What advice would you give to anyone considering starting an Open Source project?

PC: Just do it! I can't think of any reason not to. If you have the motivation to do it for yourself, then there's likely to be someone else who can benefit from it. It's much easier to get an Open Source project running than if you were trying to sell something. The end reward is probably ultimately greater. [LXF](#)

LEFT: Managing your *ZoneMinder* system is done in straightforward console windows like this.

“I find it's always easier to invest a bit of time to start with and cut down on the rework later on – the opposite of the current extreme computing trend.”

Selecting an area of the image to monitor for activity is great for protecting vulnerable areas like dark corners, fences and gates.



PHP 5



PHP VERSION 5

ULTIMATE PHP

The last 4.x.x iteration before the full-number step has just been released. Zeev Suraski outlines what great features lurk just over the horizon...

Users who consider themselves members of the PHP community, and haven't been living under a rock in the past few months, should have already heard about the upcoming version 5. However, a recent survey conducted among visitors of www.php.net, <http://zend.com> and other sites suggest that more than half of the PHP developers don't know what the new features of this version are going to be. The purpose of this article is to give you a taste of the new functionality that will be offered by the next generation PHP, and how it is going to change the way you use PHP. It's important to note that PHP 5 is no longer at 'vaporware' stage – at time of writing, you can already download Beta 1 from www.php.net, with Beta 2 coming in the very near future.

In addition to the minor changes that are customary between any release versions of PHP, the key changes between version 4 and 5 can be categorised as either changes to

the language syntax and behaviour, or major changes in the extension modules that connect PHP to external resources such as XML or SQL. In the first category, the most far-reaching change is that of the object model. PHP 5, based on the *Zend Engine 2*, features a complete rewrite of the PHP object model, and puts it on par with most other object-oriented (OO) languages. Design patterns and complex object networks, that until now were virtually impossible to develop with PHP, can now be easily implemented, with more and more examples popping on the web every day. In the second category, we can find two revolutionary modules – *SimpleXML* by Sterling Hughes, and *SQLite* by Wez Furlong, Tal Peer and Marcus Boerger. *SimpleXML* reduces the complexity of working with XML to the degree of trivial, whereas *SQLite* is going to bring the benefits of relational databases to each and every developer, whether he has access to his web server or not.

If none of this sounds interesting to you, there are doctors who can handle this problem!

Changes to the Language Syntax

PHP 5 was initially born out of one key idea – radically changing (read: fixing) PHP's object model. PHP 4's object model, which was almost identical to that of PHP 3's object model, was born out of wedlock. In the beginning, objects were not all that different from arrays, and that too is an overstatement. Arrays and objects behaved in pretty much the same way, and were actually interchangeable.

This fact, combined with the fact that the object oriented syntax attracted a huge number of users, created a problem – plenty of people using a feature that was introduced to the language with very little thought being put into it. The problem wasn't only theoretical – it was, and still is, very real. Objects in PHP 4 behave in a way that is counter-intuitive, and

very annoying to work with. Here's an example that I often use to illustrate the annoying nature of the old PHP object model:

```
<?php
function wed($bride, $groom)
{
    if ($bride->setHusband($groom)
        && $groom->setWife($bride)) {
        return true;
    } else {
        return false;
    }
}

wed($joanne, $joe);
print areMarried($joanne, $joe);
?>
```

If you try to run this code (you would have to implement the **bride** and **groom** classes, as well as the **areMarried()** function), you will not get the expected result. – **joe** and **joanne** end up being divorced just as soon as they got married. If you are familiar with the quirks of the old PHP object model, you know the reason why – the **wed()** function works on 'clones', copies of **\$joe** and **\$joanne**. These clones really end up getting married within the scope of **wed()**,



but as soon as **wed()** returns, these clones, which are local variables, get destroyed. The eternal vows of holy matrimony don't last for more than a few cycles. Technically speaking, the original **\$joe** and **\$joanne** are not even divorced – they remain single all the time.

While there's a way of fixing that in PHP 4 by explicitly telling PHP to pass arguments by reference, and in other cases – return the return value by reference, this methodology is not very intuitive nor is it easy to use. It involves adding many ampersands, cluttering the code with obscure symbols, and annoyingly, it's typically not that easy to figure out where exactly these ampersands need to be deployed. Worse yet, this behaviour, that made PHP spawn more and more copies of each object as it was passed from one function to another, made it impossible for us to implement additional OO features, such as destructors and immediate dereferencing of function return values (eg **getObject()->method()**).

With the object oriented syntax of PHP becoming more and more popular, and with the size and complexity of PHP projects reaching new heights, it was clear we had to address this issue. However, we had to find a way to do this without changing the behaviour of any of the other types, as we didn't want to affect the regular, structured code. The solution came from Andi Gutmans. He came up with an idea about how to modify the engine so that objects will be automatically passed by reference, in

a very clean and thorough way. His idea was to add a level of indirection, and have the language work with object handles, instead of directly with objects. This idea was the starting point of the *Zend Engine 2*, and in turn, PHP 5.

In PHP 5, the previous example works perfectly well, without requiring any ampersands. Objects feature an automatic 'pass-by-reference' behaviour, so Joe and Joanne no longer get away from their vows that quickly. That alone makes PHP 5 a much more friendly language to work with, when you're dealing with object oriented code. Of course, since it's sometimes desirable to operate on a copy, you can still create clones by using the special new **__clone()** method.

New features

In addition to this change in the object model, there are many more new important object-oriented features supported by PHP 5:

■ **Access level restrictions** – aka PPP (Private, Protected, Public) – it is now possible to restrict access to properties and methods so that only class methods can use them (Private), only the class and its derivatives (Protected) or everyone (Public). The default access level in PHP is public.

■ **Static modifier** – static methods and properties are now supported, allowing users to encapsulate regular functions / properties, that don't belong to any particular object instance, in a class.

■ **Interfaces** – PHP now supports interfaces, allowing for a more powerful class hierarchies and emulation of multiple inheritance (we decided against implementing 'real' multiple inheritance, because of various problems that come bundled with this feature).

A class can still extend only one parent class, but it may implement any number of interfaces. If a class declares that it implements an interface, PHP conducts checks to ensure that it does indeed implement all of the methods in the interface.

■ **Type Hints** – functions and methods can now declare what type of objects they're designed to accept, by optionally stating the class name on the left of the argument name. For example, function **Hire(Employee**

■ **Destructors** – one of the most requested OO feature in the history of PHP. The introduction of the new object model made the concept of destructors compatible with PHP. When PHP destroys an object, it checks whether a special method named **__destruct()** exists. If it does – it is called. That allows you to perform cleanup or special operations right before the object is deleted.

■ **Just-in-time loading of classes** – PHP now allows users to load their classes 'just in time', right before PHP tries to instantiate their class. If PHP tries to instantiate a class that doesn't exist, it first checks whether a special global **__autoload()** function was declared. If it finds it – it calls it, providing the class name as an argument. **__autoload()** is then



Zeev Suraski

“PHP 5 was initially born out of one key idea – radically changing (read: fixing) PHP's object model.”

Semmployee). Contrary to the belief expressed by significant numbers of users on the Internet, this is not a step on the way to make PHP a strictly typed language. This feature is designed to simplify checks for specific object types, but it will not handle the basic PHP types (integer, string, etc) that remain completely interchangeable.

responsible for loading the class into memory (typically, by calling **include()**), and once **__autocall()** returns, PHP tries to instantiate the class again. Assuming **__autocall()** properly loads the class into memory, the object instantiation goes through transparently. This feature, in addition to simplifying and reducing the number of **include()**s, can actually



PHP 5



« improve performance, as only the classes that are actually necessary are parsed and loaded into memory.

There are also several smaller object oriented features introduced in PHP 5. You can find the full list of them at www.php.net/zend-engine-2.php

“PHP used to lack a reliable storage mechanism. This era is to end with PHP5, thanks to the *SQLite* extension.”

At this point, you may be asking yourself how exactly you can utilise all of these features. While it's beyond the scope of this article to thoroughly demonstrate each and every new feature, here's a quick example that demonstrates the use of access levels, static members and a design pattern (Singleton), all in one.

```
<?php
class AppConfig {
    private static $instance = null;
    private $configDirectives = array();

    // The constructor is declared
    private so that this class cannot be
    instantiated using new
    private function __construct()
    {
        // Initialize, load configuration
        directives into $this->configDirectives
    }

    public instance()
    {
        if (self::$instance == null) {
            self::$instance = new
            AppConfig();
        }
        return self::$instance;
    }

    public getConfigEntry($name)
    {
        return $this-
        >configDirectives[$name];
    }
}
```

```
}
};
?>

Now, whenever we want to utilise
this class, all we have to do is call
AppConfig::instance()-
>getConfigEntry("foo").
```

There are several great things

about this class. A single configuration object becomes 'magically' shared across our entire application, without us having to send it throughout lengthy call paths, or having to create multiple instances of the same configuration object, and perform redundant loads of the same configuration file. Another good thing is that we don't have to pay the price that is attached to the 'customary' way of doing this, using global variables. We know that our instance is safe, nobody can overwrite it or break it in any way.

What about some non-OO stuff?

It's no secret that plenty of PHP users could hardly care less about the new object-oriented features, and stick with the structured programming style. If you happen to be one of them, PHP 5 still offers a great deal for you.

As far as the language syntax goes, there aren't too many revolutionary changes. In fact, there are only two things I can think of. One is that by-reference arguments can now have default values. For instance, you can have something like: `function foo(&$arg = null) { ... }` The other is somewhat more impressive, and rated high on the most requested features list – **foreach()** that works with references. Here's a quick example to be going on with:

```
<?php
$veggies = array("carrots",
"eggplants", "tomatoes");
foreach ($veggies as &$veggie) { //
notice the & reference mark
    $veggie[0] = strtoupper($veggie
[0]); // because $veggie is a reference
// to the array element, any
changes to it
// persist in the array
}
?>
```

However, neither of these features is groundbreaking. If you want to look for non OO features that will truly blow you away (whether you're an OO fan or not), you have to look elsewhere – in the extension modules.

SimpleXML – XML just got trivial

While not impossible, it's truly difficult to be a developer in this day and age, and ignore XML. With Web Services, RSS feeds, XSLT and various other buzzwords and three- or four-letter acronyms, we all spend more and more time working with XML. There are two key ways of working with XML – one is using an API where you provide callbacks to the XML parsing engine, that get called each time the parser recognises a certain structure. The other is known as DOM (Document Object Model), where an XML document is loaded in its entirety into memory, and you can use API functions to access the elements in the XML document, in a similar way to accessing a Tree data structure.

To make a long story short, neither of these APIs are very simple to work with. Callback-oriented programming is counter-intuitive, and requires that you do lots of bookkeeping (as callbacks are invoked for every step of the parsing, each of them provides just a small piece of additional information, such as the tag name, directive names, directive values, etc.). DOM is a bit easier to use since it provides you with a tree abstraction,

but you still have to use clumsy and long named API functions to access each element and traverse the various nodes. For instance, if we want to parse the following file:

```
<veggies>
<veggie>
  <name>Carrot</name>
  <desc>an orangey root</desc>
</veggie>
<veggie>
  <name>Tomato</name>
  <desc>a red fruit</desc>
</veggie>
</veggies>
```

and want to turn it to:

```
Carrot is an orangey root
Tomato is a red fruit
```

In DOM, a common way to do it would be this:

```
<?php
$dom = domxml_open_file("./
veggies.xml") or die("failed!");
$veggies = $dom->get_elements_
by_tagname("veggie");
foreach($veggies as $veggie) {
    $name_node = $veggie->get_
elements_by_tagname("name");
    $name = $name_node[0]-
>get_content();
    $desc_node = $veggie->get_
elements_by_tagname("desc");
    $desc = $desc_node[0]->get_
content();
    print "$name is $desc\n";
}
?>
```

If you're happy with this, *SimpleXML* is not for you. However, if you have human DNA in your cells, you probably think to yourself "why the heck should it be that difficult?" And frankly, there's really no reason for things to be that way, at least not in a dynamic language such as PHP. If you could treat an XML document as if it was a native object – with direct mapping between the object properties and the XML elements, things would become much simpler.

And, since PHP 5 is in the wish-

granting business, that's exactly how things work in version 5. Here's what the code would look like using the *SimpleXML* extension:

```
<?php
$veggies = simplexml_load_file
('veggies.xml');
foreach ($veggies->veggie as
$veggie) {
    print "$veggie->name is $veggie-
>desc\n";
}
?>
```

When you create the *SimpleXML* object, it becomes a direct mapping of the XML file. You can iterate on collections using *foreach()*, you can de-reference elements using their tag names as properties – to make a long story short, you can treat the XML object as the trivial representation of your XML file.

Compare the *SimpleXML* code with the DOM code, and you should get an immediate warm fuzzy feeling about the years to come.

Going over elements is not the only thing you can easily do with *SimpleXML*. Updating elements is equally trivial. Let's say we want to change the **Carrot** in our example, to an **Orange**. Here's how:

```
<?php
$veggies = simplexml_load_file
('veggies.xml');
foreach ($veggies->veggie as
$veggie) {
    if (strstr($veggie->desc, "orange")) {
        $veggie->name = "Orange";
        $veggie->desc = "An orangey
fruit";
    }
}
$veggies->to_xml_file('./updated_
veggies.xml');
?>
```

In the spirit of PHP, *SimpleXML* doesn't settle for providing access for the simple stuff only, but provides full coverage of the features of *libxml2*, including searching using *xpath*

(**SimpleXML::xpath()**), validating schemas and more.

To summarise, with *SimpleXML*, PHP 5's XML support can be rated Excellent, and in a world where the usage of XML for data storage and communications is ever growing, that's a big advantage.

SQLite – goodbye to custom data storage!

One of the problems that many PHP application developers bumped into at some point or another, is the lack of availability of a reliable storage mechanism. The Catch-22 was simple – if you wanted a data store that was reliable (namely, an SQL backend), you couldn't rely on it being on your target machine. And vice versa – if you wanted a data store that will surely be available on all target platforms, you had to revert to an unreliable data store – custom files.

This era is about to end with PHP 5, thanks to the *SQLite* extension, that will be bundled inside PHP. Since remarkably and for no good reason in my opinion, *SQLite* has gone unnoticed for years and years, it may require an introduction.

SQLite is an SQL library, which implements most of ANSI SQL 92. Note the word 'library' in there, and more importantly, the lack of the word 'server'. *SQLite* is not an SQL server, but an embeddable library that implements SQL. When you work with *SQLite*, there is no server. Instead, the library performs as both the client and the server, and translates the SQL queries directly to operations on the local database files. For small and even medium scale applications, this makes perfect sense. Why bother with client/server architecture, when all you want to do is save some data in an easy-to-use format?

At the top of the next column is a simple script that utilises *SQLite*:

```
<?php
$conn =
```

```
sqlite_open("company.sqlite")
or die("Unable to open database");
$query = sqlite_query("select name,
title from employees") or die("Query
failed");
while (list($name, $title)=
sqlite_fetch_array($query)) {
    print "$name's title is $title\n";
}
?>
```

So far, it looks simple enough, very similar to other SQL extensions that PHP supports. But it gets better than that. Consider the following example:

```
<?php
function noChiefs($title)
{
    if (strlen($title)==3 && $title[0]
=='C' && $title[2]=='O') {
        // One of these CXO people...
        return 0;
    } else {
        return 1;
    }
}
```

```
$conn =
sqlite_open("company.sqlite") or
die("Unable to open database");
sqlite_create_function($conn,
"NO_CHIEFS", "noChiefs");
$query = sqlite_query("select name,
title from employees where
NO_CHIEFS(title, $conn)") or
die("Query failed");
while (list($name, $title)=
sqlite_fetch_array($query)) {
    print "$name's title is $title\n";
}
?>
```

This example illustrates the interesting nature of *SQLite* – since it is running at the same place as PHP, you can make use of PHP functions inside your queries! In this example, we used it to easily get rid of CEOs/CTOs /CFOs in our query. In order to do that, we first need to register our PHP function with *SQLite* and give it a name that *SQLite* would recognise, using **sqlite_create_function()**. Afterwards,

we can simply use this function in our *SQLite* queries.

Another interesting aspect of *SQLite* – that makes it an even more natural partner to PHP – is the fact that it is typeless. Like in PHP, there's no need to worry about the type of the data, the only thing that matters is the value.

The name of *SQLite* can be considered to be a bit misleading – other than the fact that its footprint is small and that it doesn't come with the bloat typically involved with an SQL server, there's really nothing 'lite' about it. It's fast (typically faster than MySQL and PostgreSQL, see www.sqlite.org/speed.html) and it is full-featured, including functionality that was only recently introduced to some of the bigger Open Source database backends such as transactions, views, and more.

Each of the above-mentioned features alone is extremely useful. Combining all of them, and adding the fact that *SQLite* is bundled in PHP 5 and is enabled by default, and you quickly realise that it's revolutionary.

Conclusion

PHP 5 is a big and important improvement to the already impressive and powerful PHP platform. Whether you're interested in object-oriented development, XML development or reliable data storage, PHP 5 is bound to radically improve the way you do things on the web. Coming soon to a web server near you! [LXF](#)

ZEEV SURASKI

One of the designers of the engine behind the Open Source programming language PHP, and a member of the Apache Software Foundation, Zeev is also CTO of Zend Technologies, a company that develops products for the PHP market.



What on Earth is... “THE GRID?”

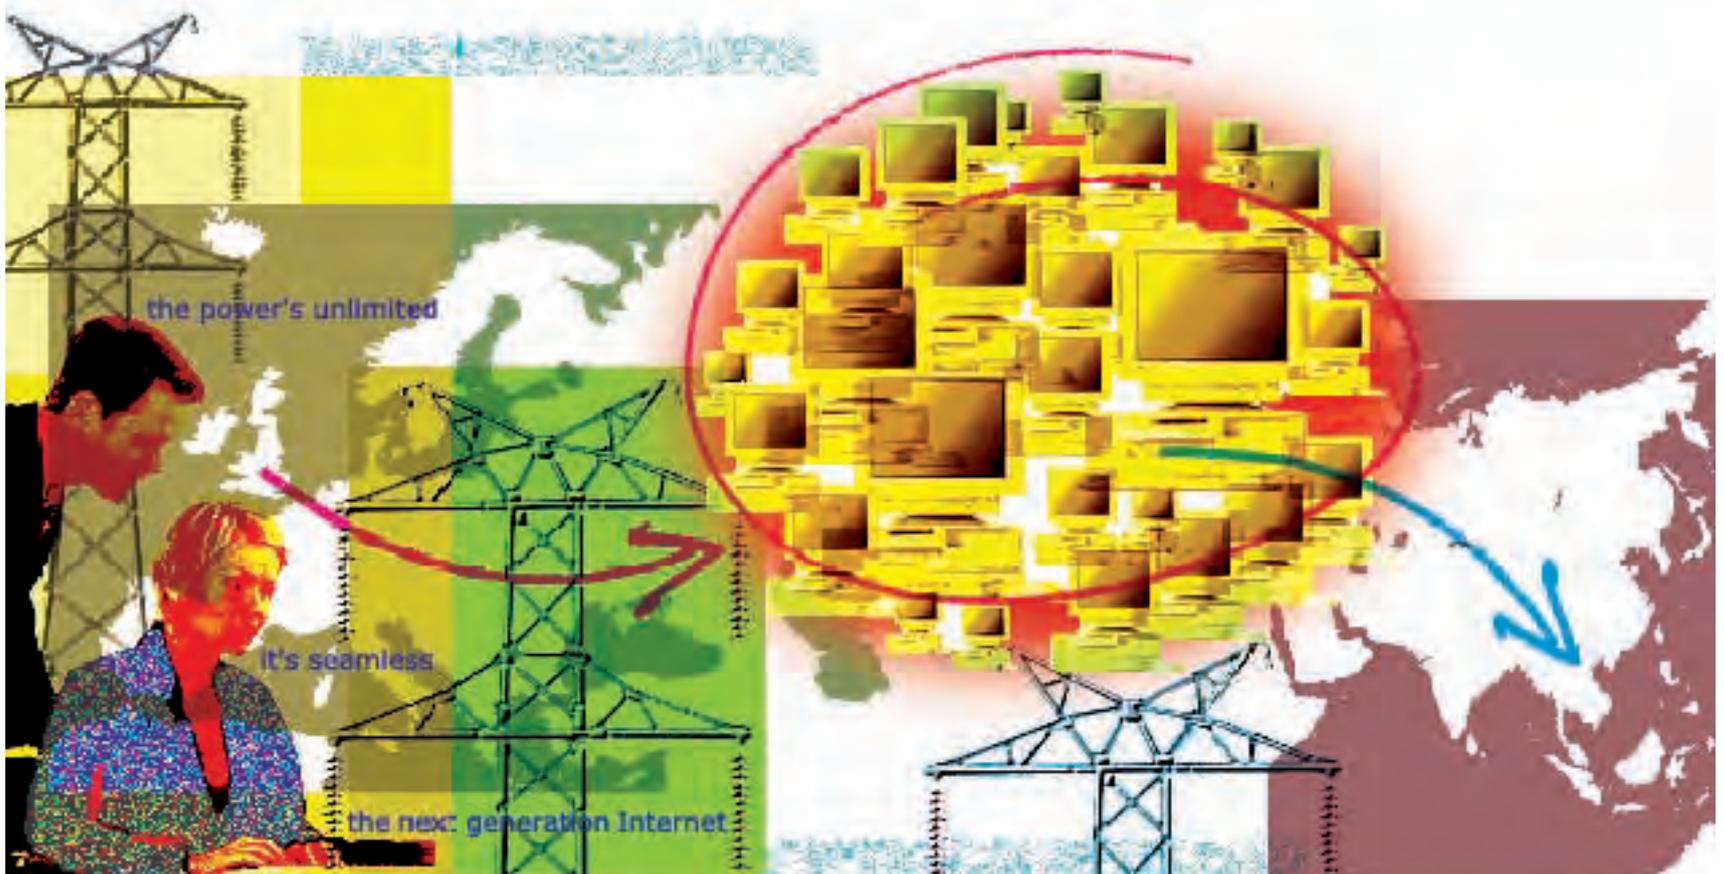
There will come a point when almost every PC on Earth is hooked up to the Internet. But what will happen then? **Andy Channelle** probes one of the most exciting concepts that is happening to computing...

» Can we get the wild claims and hyperbole out of the way first? What is the Grid then?

OK.. Hitch a ride back, if you will, to the year 1980. You are talking with a knowledgeable computer user who suggests to you that a small application, called Enquire-Within-Upon-Everything, created by a British scientist in Switzerland will over the next 20

years completely revolutionise information management and retrieval. We speed forward 10 years where Tim Berners Lee – for it is he – comes up with a proposal (variously called *Information Mesh*, *Information Mine* and *WorldWideWeb*) for a global hypertext network which gets the go ahead from his bosses at CERN. This, with a few graphical refinements, is the web we know and love.

The Grid, then, could be called the next evolutionary step for the Internet; the linking of the ‘traditional’ net with peer-to-peer and distributed computing technologies. Some scientists go so far as telling us that the Grid has the capacity to change everything we do – from jotting down shopping lists to crafting high-resolution photo-realistic film special effects; from e-commerce to home automation.



» Wasn't that the ultimate goal of the Internet in the first place?

Not really. The only goal the originators of the Internet, and later the world wide web, had was making documents and research papers available quickly and simply to (both military and civilian) scientists around the world. Everything else has grown out of that. In fact, one could argue that the Grid is the result of these scientists finally turning their big brains to the potential of the Internet, which until now has been driven – to a large extent – by the needs and ambitions of artists, software developers and commerce.

» Oooh, I should expect some flames about that statement! How is the Grid different from the World Wide Web and other stuff I use?

When you sit at your PC and type a web address into *Firebird* or *Konqueror*, a document is retrieved via the Hypertext Transport Protocol (HTTP) from a server. This could be from a vast server farm with 8,000 nodes, or it could be from a wheezy 486 sitting on the end of a dial-up connection. While these are vastly different departure points for a document, its destination is the same: your PC with its 1GHz processor, 256MB RAM and 40GB hard disk – or whatever specs your system has.

The situation with peer-to-peer is different, as in addition to being a client of the system, your own computer also becomes a server providing access to documents, MP3s, software etc to other users on the network. Think of KaZaA, Gnutella and BitTorrent. What stays the same is your PC with its defined and static specification.

Distributed computing (*What On Earth, LXF42*), meanwhile, parcels off bits of data to machines – or nodes – on the network for analysis or processing. Prime examples include SETI@Home and the 'farms' that rendered scenes from *Shrek*. And while your PC may become – if only for a few processor cycles – part of a bigger entity, it is still the same machine with the same processor and the same limitations.

» How is Grid computing any different from these?

According to IBM's Grid Guru, Brian Carpenter, the Grid will bring about the commoditisation of computer power, putting processing resources on the same footing as any other 'utility'; gas, electricity, water. He told *Wired* magazine that the Grid will "open up storage and transaction power in the same way the Web opened up content"

The big idea is that the Grid will eventually be accessible from any compatible device – desktop PC, set top box (STB), PDA – giving access to a vast amount of processing and storage power regardless of the status of your client. The network is pervasive in the same way as the electricity grid – you don't have to worry that your fridge isn't going to be compatible with the electric sockets in your house, it

should be the same with your computer/device. Ian Foster and Carl Kesselman, in their seminal work, *The Grid: Blueprint for a New Computing Infrastructure* (Morgan Kaufmann; ISBN: 1-5586-0475-8), described the Grid as "a hardware and software infrastructure that provides dependable, consistent, pervasive and inexpensive access to high-end computational capabilities"

As the difficult bits (processing, data storage, etc.) are removed from the local device, size and initial costs can be reduced, while battery life for mobile devices can be increased. For example, a journalist may be taking notes on her low-spec handheld PC at a news conference, but the (rented) application she is running is on a server in Tokyo and the data may be stored on the company server in New York. And all this happens as though both application and data are local.

» Sounds cool: a super computer in your pocket! But I thought the trend was toward decentralised PCs of increasing complexity?

It has been in the past (and that suits a number of companies), but there are very good reasons, especially in the area of high-performance and scientific computing to distribute or share resources. There are also potential benefits at the other end of the scale too.

For 'low-end' users who may do a bit of web grazing, send emails, write a letter and not much else, there's little point in splashing out on a Pentium 4 with hyperthreading, 1GB of DDR RAM and a vast hard disk. There are less expensive ways of keeping your right leg warm!

This user could have a 'stateless' terminal – screen, keyboard, mouse – which connects to the Grid on demand and allocates the resources required to do a specific task. This device could be rented just as you rent a cable set-top box, with a basic set of services and storage options. As his or her needs grow, the system scales, so if our user one day decides to do a little oil prospecting – which would obviously necessitate some high-resolution geophysical data processing – more Grid resources are called upon without the need to completely re-equip the living room. Inevitably, as need scales, so does cost.

At the other end, a particle physicist in Switzerland could organise it so that, for a short period, he could monopolise the Grid and bring the power of 8,000+ processors to bear on analysing his data. Moreover, the ultimate goal is that, when required, disparate Grids could be lashed together seamlessly to create what scientists have named a metacomputer...

» ...which could then take over the world and enslave humanity...

Don't be silly! The whole concept does sounds a little sci-fi (and utopian) but in Japan a consortium of businesses, universities and government departments



WHAT ON EARTH The Grid

« has begun a project which will link together thousands of computers, including existing supercomputers, clusters and even unused desktop machines using Grid principles to create a beast rated – if need be – at 100 teraflops. The most powerful single computer in the world at the moment, NEC's Earth Simulator, is rated at 36 teraflops.

That's amazing in itself, but the Naregi (National Research Grid Initiative) plan aims to be compliant with the Open Grid Services Architecture (OGSA), meaning it could become just one more node in an even bigger Grid encircling the globe.

No-one mention White Mice or Slartibartfast! One of the most famous quotations in the history of Information Technology is Chairman of IBM Thomas Watson's seemingly laughable assertion back in 1943 that there was a worldwide market for perhaps five computers. Grid acolytes would suggest that Watson overestimated by four!

» Isn't this all a rehash of the 'The Network is the Computer' stuff that Sun was banging on about all those years ago?

To a degree it is similar – in that local resources are broadly irrelevant, and it has to be said that the concept of Grid computing is being warmly embraced by those corporations that have a large stake in providing 'back-end' services (IBM, Oracle, Sun etc.) or infrastructure.

In the past though, the involvement of so many corporate vested interests led to conflict over standards, which is why the scientific community building on the foundations of the Web – which has resisted many attempts at creating non-standard standards – is making headway while hardware vendors pull in different directions to achieve the same objective.

Things are changing though. For the Grid to work, applications created by Oracle must work with middleware from IBM on hardware developed by Sun over infrastructure built by Cisco: hence the notion of Linux – with its dedication to openness, low cost and network transparency – as being the ideal Grid operating system.

» I've played a few online games and often the latency causes problems (and this is where most of the processing is going on locally!). There's no way the Internet can handle this much data at sufficient speed, is there?

That may be true at the moment, but Grid experts claim to have learned the lessons of the Internet, noting where and why bottlenecks have occurred. So the next generation of public networks will be built to take these in to account. In fact, the cable networks built at enormous expense by the likes of NTL and Telewest should have the capacity to run Grid services in the future. And don't discount the advances taking place in the wireless arena: Ethernet speed over distances of 30–40km is not out of reach with current technology, so improvements in this area have the potential to eliminate the 'last mile' bottleneck that some claim is holding up broadband technology in many parts of the UK.

» Sounds like a lovely dream. How long before we have some workable Grid applications?

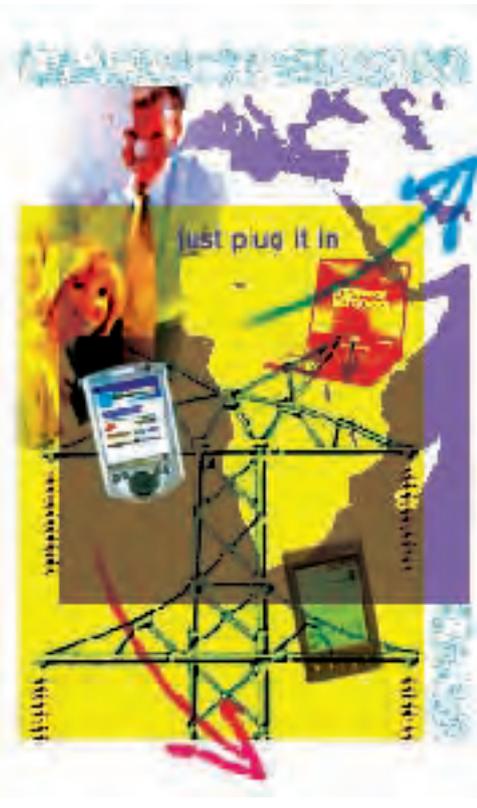
They are already here, sort of. Grid applications will probably follow the course set by the Internet, which means the technology will be used by – and perfected – by the scientific community, perhaps for years, before making its way out of the lab and in to the living room – though Sony may prove me wrong in early 2005.

Between science and leisure is enterprise, and there are already a number of established companies and startups pushing the benefits of Grids to business. IBM, for instance, has recently announced a new version of *Websphere* which can use Grid principles to hook up every PC and mainframe in an organisation to create a single virtual system which allocates resources as they're needed. There are limits to the capabilities of *Websphere* at present – it only works with a single cluster on a specific application – but future versions should allow the linking of multiple clusters throughout an organisation and run many disparate applications.

It's all moving fairly rapidly. Don't forget the Internet was a geek's plaything less than a decade ago, and Grid developers are not keen on wasting so much time before going mainstream with 'Internet version 2.0'.

» I'm inspired! How do I get a piece of the Grid action?

There are a number of public projects using Grid technology today even if they are not going to set

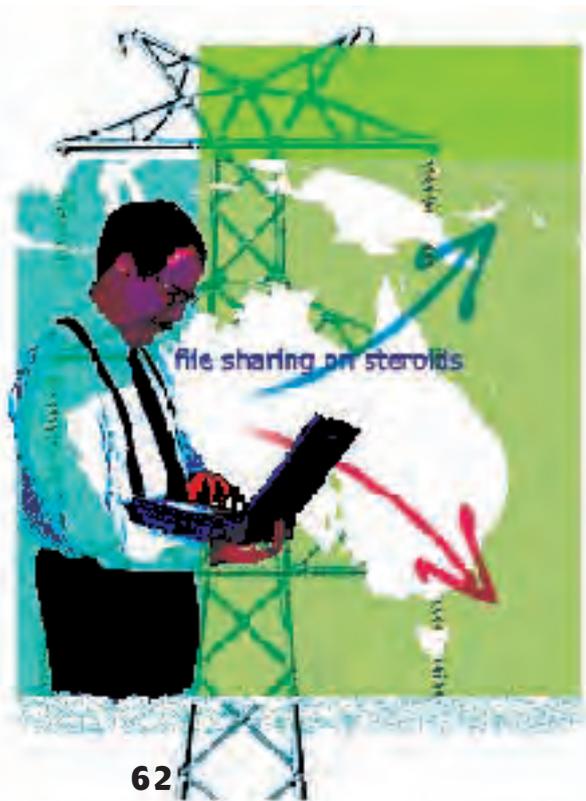


the world of consumer computing alight. In fact, you're probably helping to fund one of the more important projects via the EU, so why not have a look at it?

The EU DataGrid project was started by a number of organisations including the European Space Agency (ESA), CERN and France's CNRS, and has a suite of three applications in development, covering high-energy particle physics, genomics and Earth observation processing. These three fields commonly grow up around what DataGrid calls Virtual Organisations (VO), that is a community of several institutions and individuals "sharing the same interests and, to a certain extent, the same goals." Until recently these communities may have had the same goals, but sharing resources in a 'transparent, coordinated and secure way' had been all but impossible, not least because of a blurring of boundaries between academic and enterprise spheres.

The current state of the art involves a four-stage process – submission, search/allocation, monitoring, result – and, as you'll discover, is not at the point where you can simply load up a remote copy of OpenOffice.org and start typing.

- The first stage is to submit a request. As well as including the input data, this would specify the high level requirements of the job: operating system, application type and maybe even a named application. The EU DataGrid project uses a web-based UI for this stage which, in the case of the ESA Earth Observation Grid, is indistinguishable from a standard web form.
- The Grid then searches for and then allocates the required resources, including processing power and storage requirements, and begins the job at hand.



- As the job is being done, the Grid will monitor request processing in real time so that resources can be dynamically scaled depending on requirements.
 - Finally, the Grid notifies the user once the job is complete and presents its output.
- As I say, it's not something the average user is going to get excited about yet, but the potential power available, once the world is Grid-enabled, is staggering. As an example, ESA's Earth Observation testbed Project involves collecting some 100GB of raw image data per day. From this, the Grid extracts specific atmospheric ozone data for analysis. Although this is useful in itself, the ESA Grid's main purpose is to demonstrate the improvements in accessing and processing enormous datasets stored in a continent-wide distributed network.

» Good news for science and computing, I suppose. But what about me? But did I read something in the Web about the Playstation 3 being built on Grid principles?

Hmmm, sort of. The Playstation 3 is going to use a group of 'cell' processors – reports vary between four and 16 processor cores – and Sony says the machine could, in theory, form the basis of a Grid, just as the PS2 has been used to create a Linux cluster. That said, there are a number of problems to be overcome, the biggest of which would be convincing PS3 users that donating spare resources to strangers is a good idea. However, Grid principles would appear to be the only way for Sony to bring about the promised 1,000-fold increase in power for its next console, and may form the basis for a whole range of new consumer goods.

There are interesting applications of Sony's technology: your cell-based television – which may only house one processor core – could combine with your cell-based DVD player (two cores) to help

out when your Playstation3 (16 cores) begins struggling to render real-time photo-realistic scenes from Final Fantasy XIV. How about a Grid for each and every home?

There has also been an attempt recently to put Grid principles to demonstrable use in gaming using the Open Source version of *Quake II* and some IBM middleware. *GameGrid* was developed by students at the University of Wisconsin on the foundations of IBM's *OptimalGrid* product.

Most *Quake* games, an IBM spokesman said, would push the server with 32 players; *GameGrid* (built with eight servers) served up a specially designed map to 80 players without breaking a sweat and could have taken many more. It also provided *Linux Format* with a simple way of breaking down the Grid concept!

GameGrid's game world is divided into sectors, each one housed on a different server. As a player or object moves from one sector to another, they also switch servers. A rocket, for instance, may shoot down a corridor from one server into a room on another, what passes between servers is statistics – object name, velocity, vector *etc* – with a typical latency of less than 50 microseconds.

Where the Grid really comes into its own though is when 80 players all converge on one sector of the map, potentially overloading that sector's server. In this case, *GameGrid* dynamically allocates resources from the rest of the Grid to balance the load. Moreover, future Grid clients could rope local resources (where they exist) into the system, so the more people who play, the less strain; just as the popularity of a BitTorrent file reduces each users download time.

» It still all sound very similar to the distributed computing efforts I've read about.

It's a similar concept but in the case of *GameGrid*, it's important to remember that the 80 users playing on the system, though some of them may have been playing on top of the range PCs, were using their local systems merely to display and control the game, all the processing was going on at the back end. Playing on a connected PDA (with a decent screen) would not have disadvantaged the user.

» Is there really a need for this much computer power, and how will it make a difference to the average person in the street?

In addition to the scientific stuff that seems a little remote, and games that may be irrelevant to most people, there are a few really astonishing projects under way which will have a fundamental impact on our lives over the next few years. One is *Mammogrid*, which aims to improve breast cancer survival rates and, as an offshoot, develop a secure model for remote database access which may eventually benefit the entire UK health service.

Every year, 1.5 million women between the ages of 50 and 64 are offered a mammogram. An initiative which annually generates three million images to be analysed. Each image is viewed first by one of the UK's 230 trained radiologists and suspect cases reviewed by a senior radiologist. Anomalies are found in 30 cases per thousand, yet 20 per cent of cancers are missed.

eDiamond's *Mammogrid* project will use the extraordinary data handling facilities of Grid technology to manage and analyse images (using 'find-similar' principles to create a tumour recognition system), and also take care of the data and security features (authenticated, verifiable access *etc*) that the law requires.

The project's leaders are looking at a variety of Grid structures including a system where patient records can be held in a central database, while mammogram images can stay at the patient's local hospital, yet still be available for comparative purposes to the rest of the network. Alternatively, a hospital could retain 'ownership' of its data with the central server hosting a federated view of the data. A query would be farmed out across the Grid and the results collated before being presented. In both cases, the network manages storage and allocates resources to the necessary analysis.

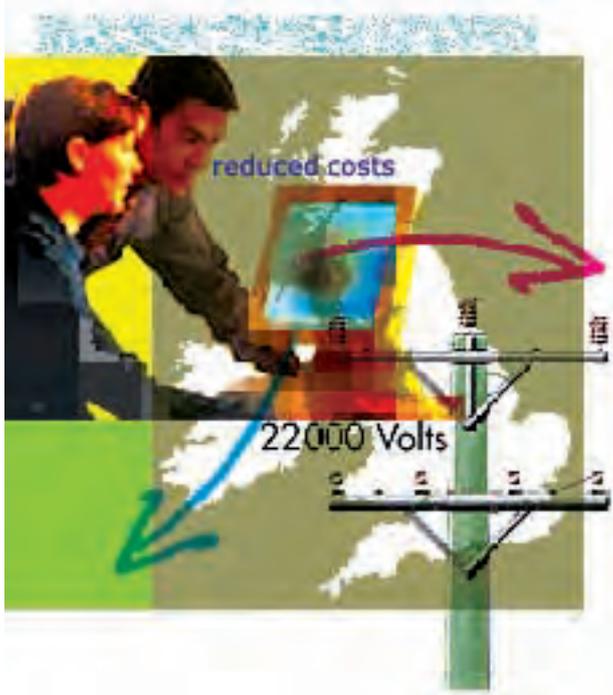
Another life saving solution being planned at the National eScience Centre in Edinburgh, is a fire modelling system which can take real-time data from fire crews attending a blaze and sensors in the building and tie it up with stored architectural plans to predict the spread of the flames. This could offer firefighters the best route through a burning building.

Both of these applications require not just vast stores of data, but rapid but sporadic access to massive processing power, the things that Grids are especially good at. It will also require a *lot* of investment, but many pundits believe that, having been caught on the backfoot with the Internet, the authorities will be keen to be seen exploiting such cutting-edge technology.

» What with this being a Linux magazine, it would be rude of me not to ask what role the Linux operating system plays in all of this.

Linux is actually in a unique position to become the *de facto* Grid standard. Not only does it power the majority of Internet v1, it is also central to the future plans of the hardware and software companies that will push Grid computing forward. Look back at announcements of supercomputers over the past couple of years and you'll find that almost all of them are being designed around what was, just a few years ago, regarded as a hobbyist project.

The Grid will further cement Linux's reputation in enterprise and may well even drive desktop use as CFOs realise that 10,000 PCs running screensavers all through the night could surely be put to much more profitable use. 



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

Got a problem with your browser rejecting sites that are *IE*-only? Here's how to spoof them... **p66**

Regular Expressions

It's hard to overestimate the importance of regular expressions in Unix/Linux. They crop up all over the place! **p70**

Compiler writing

How to go about extending SKYLang so that it handles conditional statements **p74**

Blender

If your scene needs fog, swarms of insects or realistic fire effects, a Particle system is the answer **p78**

The GIMP >>

Photo manipulation – how to restore your treasured photographs to their former glory **p82**



PHP

Spread your bandwidth more thinly – output buffering can compress content before you send it **p86**

Server School

Everyone should have a Linux server! All you need to know about getting to grips with *Apache 2* **p90**

TIP OF THE MONTH!

SysRq MAGIC

Long considered perhaps the most useless key, you might be surprised to hear that the System Request key on your keyboard is actually remarkably helpful in Linux. On most standard keyboards, it's somewhere above the Insert key, but natural-style keyboards tend to have it somewhere else – often above the numeric keypad. On most (all?) keyboards, SysRq shares its position with Print Screen, another fairly unused key on our keyboards.

Anyway, onto how SysRq can be a lifesaver. If you have configured your kernel specially, or if you are using a distro that comes pre-configured to use SysRq properly such as Mandrake

9.1, then you can take advantage of the 'Magic SysRq' feature of Linux. If you press **Alt+SysRq** and then one of several keyboard keys, the Linux kernel will automatically spring to life and help you out.

For example, if you ran a bad program that locked up your PC, how do you kill it without SSHing into the box? With Magic SysRq you have two options: using **Alt+SysRq+K**, you can kill all processes on the current console, including your current login. Alternatively, to force a clean shutdown of your machine, press **Alt+SysRq+S**, **Alt+SysRq+U**, then **Alt+Sys+B**. Individually that Syncs all

open kernel buffers to disk, then unmounts the root partition and remounts it read-only, and then finally reboots.

This functionality can be called on by any user, root or otherwise, which could be a potential security problem. However, this is why it usually needs to be enabled when compiling your kernel – you need to be absolutely sure you want it enabled by enabling the option in Kernel Hacking. When using SuSE, the option is compiled into the kernel by default, but is only enabled if you edit the file `/etc/sysconfig/sysctl`, and set **enable_sysrq**, then run `suseconfig`.

GET CONNECTED

Beginners' Guide to Linux – browsing the Internet

PART 5 Without a connection to the outside world, your box is really only half a computer. Fortunately, Linux was built with networking and Internet accessibility at its very heart, explains **Andy Channelle**.

Not so long ago, having a computer connected to the Internet set you apart from the hoi-polloi. Now, it's unusual to find a PC that isn't jacked in to a phone socket, LAN or cable modem, and access to both email and the World Wide Web is held to be as vital to modern life as television and telephone.

In this tutorial we'll be covering the basics of web browsing. And while we may highlight the use of one application over another, most of this knowledge is transferable. Don't be afraid to experiment with the setting in your browser; you're unlikely to break anything!

Once you're up and running, please come and join us in the Magazine section of the forums at www.linuxformat.co.uk to discuss this and other tutorials in the series.

The Web

The first and most basic thing you'll do with a web browser is type a URL (Uniform Resource Locator), or web address in the address bar at the top of the browser window. These addresses begin with **http://** (see *Acronym Overload* box overpage), but as this is a standard thing there is no need to type it, the browser

does it automatically. The web address proper usually begins with **www.** and then the site name – **linuxformat** for instance – followed by the top-level domain (TLD) which often denotes the nationality and status of the address. For example, *Linux Format* can be reached with both **.co.uk** and **.com** suffixes; both are commercial TLDs, but the former is UK-specific.

www.greenpeace.org, on the other hand, uses a TLD usually reserved for non-profit or charitable organisations.

The revolutionary thing about the web is the ability to link documents together, so a news story on, say, the BBC's website, may link to resources around the world. Clicking the link – usually denoted by the text of the link being underlined – will take you to the new document as though you were flipping the page of a newspaper. But links don't have to be text, they could also be images designed to look like buttons or even photographs.

A well-designed site will have a consistent method for linking that is easily decoded, with links well highlighted to avoid confusion. A badly designed site may obscure its navigation leaving the user to hunt around for goodies or, as is often the case, leave and go somewhere else.

In this tutorial, we outline the best way to 'do the basics' in web browsing: bookmarking, setting up a homepage, searching, a brief look at how to convince the world you're running some other browser on some other OS (spoofing) and how to deal with one of the web's most annoying annoyances – popups.

Konqueror

BOOKMARKS When you find a site worth revisiting, it makes sense to 'bookmark' it. That is, install a small link in the browser to save having to remember and type the address. Like links and addresses, Bookmarks are a standard part of web browsing, but most browsers like to do them in their own particular way.

Konqueror tends to shove Bookmarks on a toolbar of their own. While this can be handy, without decent organisation, the toolbar soon become a mess – you can turn it on or off by right-clicking anywhere in the bar and selecting 'Toolbars>Show Bookmark Toolbar'. It is much better, however, to take some time to organise your bookmarks in such a way that the toolbar becomes your friend, with websites broken down into sensible divisions.

You can launch *Konqueror's* bookmark manager by going to Bookmarks>Edit Bookmarks. This applet works just like a normal file manager – you can move things around, create new folders or bookmarks and rename entries – and when you've finished refining, exit out and the changes are reflected in the browser. If you're a more 'spontaneous' person, you can easily add bookmark folders as you go along, by selecting 'Bookmarks>New Bookmarks Folder.'

BBCi is the king of consistent linking. Links are highlighted in blue and become underlined as you 'roll over' them.





KDE's Bookmark system works like a standard file manager.

Now when you come across a memorable site, you can click into the appropriate folder on the bookmarks toolbar and select 'Add Bookmark'. Revisiting the site then simply involves selecting the appropriate folder and bookmark. The real bonus if you are using KDE and *Konqueror* is that your bookmarks are automatically available under the Bookmarks entry in the K panel, clicking on one will open *Konqueror* at the right location.

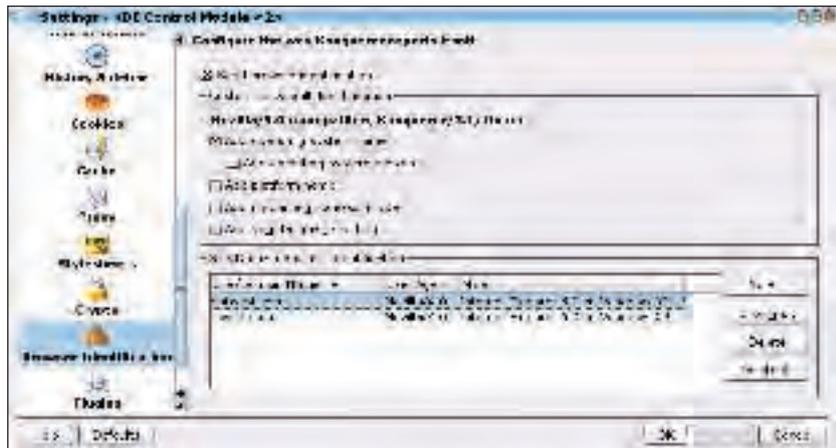
If you are migrating from another browser, *Konqueror's* Bookmarks editor allows you to import favourites/bookmarks from *Internet Explorer*, *Mozilla*, *Galeon* or *Opera* and will also export your collection of links if you're moving to a new browser. The options are accessible under 'File>Import/Export'.

HOMEPAGE By default, *Konqueror* launches with a 'help' window or, in many of the common distributions, the vendor's homepage. This can be useful the first few times you run the application, but pretty soon you'll want to change it.

You can use either a website or local file for your homepage and changing it, as with most of *Konqueror* involves going into the 'Settings>Configure Konqueror..' dialog. To use a website, simply type the address (including the **http://** part) into the Home URL bar, while choosing a file just involves clicking the small folder icon at the right edge of the bar and browsing to the file you want to use. To restore the URL to its former state, just put a tilde (~) in the address bar.

TIP

Have you noticed on most application menu bars that each entry has one letter underlined? This just means you can open that menu – or select an entry – by using the **Alt** key on the keyboard and hitting the right letter. **Alt+F** always launches the File menu, **Alt+E** launches **E**dit etc. As well as the underlined letters, some tasks also have a 'keybinding', which will be noted next to the menu entry: **Ctrl+C** copies a file or portion of text, **Ctrl+P** pastes it, and **Ctrl+Q** will usually exit (quit) an application.



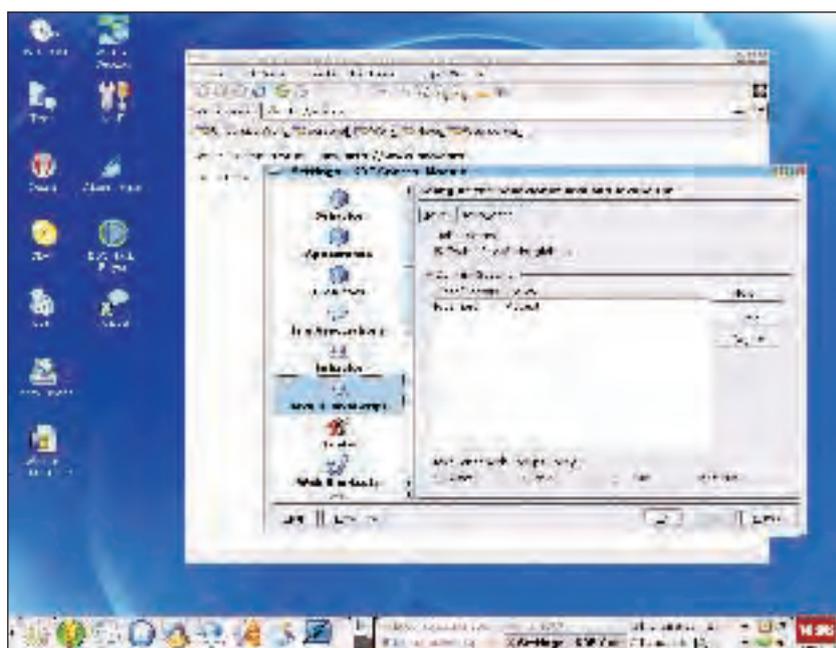
SEARCHING If you routinely use Google for your searches, *Konqueror* makes it easy to access your favourite engine, and others. Before you start, go into 'Settings>Configure Konqueror..' again and select the 'Web Shortcuts' tab on the left. Make sure 'Enable Web Shortcuts' is selected.

Back in the main browser, type **gg linux**. This should send off a Google search for Linux, and you can, of course, substitute **linux** for whatever you're searching for! You can go directly to the highest rated link for a Google search by using **gg1** as a prefix. Check out the Web Shortcuts tab to find more prefixes or, if you're feeling particularly adventurous, click the 'Add' button to create your own shortcuts.

SPOOFING Some webmasters just don't seem to understand that there are browsers out there other than *Internet Explorer* running on Windows, and there are Linux users frustrated at being denied access – often – to their own data. Banks are especially good at this, but it doesn't have to be so.

Most websites can be convinced that you are in fact using *IE* on Windows (or some other browser/OS combination) by a small manipulation of the browser's User Agent string. This used to be quite a difficult job, but now involves nothing more challenging than a little trial-and-error. For instance, we have had success convincing NatWest's online banking site to let us browse and

With a little work, *Konqueror* will get past most uninformed web designers' *Internet Explorer*-centric creations that might otherwise be unviewable.



Popup blocking will be a welcome addition if you spend significant time on the web.

TUTORIAL Beginners' Linux

◀ manage our account details by selecting 'Tools>Change Browser Identification>Internet Explorer 6 on Windows XP'; and then navigating to the site.

One problem with this is that it can distort webmasters' impression of the size of their Linux using audience, so it is a good idea to spoof just the sites that need it. To do this, do 'Settings>Configure Konqueror..' and choose the Browser Identification tab. Here you can input individual web addresses and choose how your browser is identified by them.

If you are spoofing a website, it might be worthwhile letting the webmaster know that **a)** you are annoyed at having to jump through hoops to view the pages on a non-MS Windows system, and **b)** it works beautifully in your choice of browser.

POPUVS Popup advertisements blight the lives of general browsers, and can especially cause real problems if you unintentionally stumble across one of the Internet's less salubrious quarters. You don't have to suffer – just don't accept them! Go into *Konqueror's* configuration dialog and choose the 'Java and Javascript' section and hit the Javascript tab.

The browser can be set to 'accept' or 'deny' all popups, or 'ask' you for a response every time a pop up attempts to launch. There is also a 'smart' option which tries to judge whether a popup is useful or not, and while it works in general, it is not yet smart enough to know your intentions, so use it with caution.

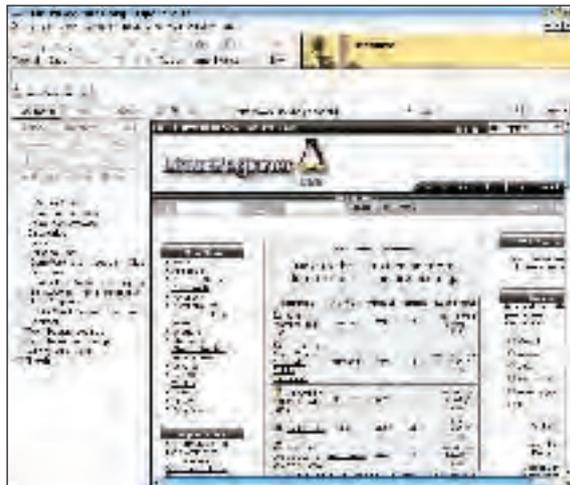
Opera

BOOKMARKS *Opera's* Bookmarks system shares a lot of ideas with *Konqueror*, including a file manager style application for organising bookmarks and the handy 'Add Page Here' command which allows you to drop links into pre-defined folders. You can't create folders from the main browser window though. The management app is accessed via 'Bookmarks>Manage Bookmarks'. In this dialog box you can add, remove, move and edit bookmarks or import them from a range of other browsers, including *Konqueror*.

In *Opera* you can access your Bookmarks via its own menu entry, or open the Hotlist (F4) and select the Bookmarks tab.



Opera has a range of configurable panels to facilitate easy customisable browsing.



Access your *Opera* bookmarks via the Hotlist by hitting F4.

HOMEPAGE *Opera* is the most configurable browser at startup time. Open 'File>Preferences' and look under the 'Start and Exit' tab. You can start up with a home page – simply add the URL to the address bar – with the page that was open when you last shut down or with a 'Saved Session' which could include a number of sites under different tabs. For the last option, you will have to set up your session: when you have it exactly as you want it, select 'File>Session>Save Session' and give it a name. You can then choose this file to start with in the Start and Exit dialog.

SEARCHING *Opera* has two embedded methods for searching: on the main address bar there is a short text box which allows searches of Google, and the Hotlist (F4) also has a search tab which is more fully featured. While Google is the default for *Opera's* search bar, just clicking on the arrow to the right of the text box reveals no less than 18 options, including Google's Linux and Image searches, Amazon and All The Web.

To add Google (or another search engine) to the Hotlist, right click in the top section of the panel and select 'Get Panel!'. The main window will show a web page with a range of panel options; choose 'Search' and then Google. You can also add news sources such as Slashdot, Wired News or BBC headlines in the same way, and each panel gets its own tab at the head of the Hotlist.

SPOOFING *Opera's* User Agent options aren't as extensive or effective those available in *Konqueror*. In fact, most websites we

WINMODEMS

Some *do* run under Linux

Linux has a great track record for supporting common and not-so-common bits of hardware, but it still has a reputation for being difficult and unsupported by the mainstream. Much of this stems from 'the Winmodem problem'.

Modems are devices that allow you to hook up to the Internet through a normal phone line and a Winmodem is an ultra-cheap (and so, popular) device that offloads much of the MODulating and DEModulating to Windows. Manufacturers like them, because they cut the cost of system building; while Microsoft likes them because they usually only work with Windows.

Usually? Yes, there are some Winmodems that work with Linux! In fact, as time passes, more and more are becoming available.

Unfortunately, getting the blighters working will involve a full tutorial in itself, but in the mean time, we'll direct you to www.linmodems.org which has some useful information and a good selection of links to help you get up and running.

ACRONYM OVERLOAD

Abbreviations explained

Many Letter Acronyms (MLAs) are not, contrary to belief, designed to make novice users feel inept, they do actually have a use! You don't need to know what they all mean, but a basic understanding will make your computer experiences a lot less surprising.

HTML HyperText Markup Language. When you visit websites, the pages you view through a browser are created with HTML. This is a relatively simple 'programming language' that uses tags, enclosed in <> brackets to format text and images.

HTTP HyperText Transfer Protocol. The method used for moving HTML code (and sometimes files) from a server to your PC where it is rendered by the browser.

FTP File Transfer Protocol. Commonly used for transferring files from server to client. There are many ftp clients available for Linux such as gFTP and kFTP, but you can also use Konqueror or Nautilus to access FTP servers.

POP Post Office Protocol. A quaint name for a common email protocol. If you're accessing a Yahoo!, MSN or other webmail account via an email client such as Evolution, you will be using POP3, which is the third version of the protocol.

IMAP Internet Message Access Protocol. An alternative to POP, which allows mail and folders to be remotely manipulated in the same as local mailboxes.

SMTP Simple Mail Transfer Protocol. A primitive mail transport system that, despite its numerous inefficiencies, is still very popular.

SSL Secure Socket Layer is the most common way that websites encrypt data between client and server. SSL is what ensures your credit card details are not intercepted by malicious parties en route from your PC to Amazon.

tested that expected IE on Windows didn't display at all, and those requiring it to access a secure server also failed. It does work occasionally, but don't rely on it.

POPUPS Opera's popup denial is very comprehensive. To change the settings (it ships with 'allow' as the default), simply do 'File>Quick Preferences>OPTION!'. We've found the best choice is to 'Open Requested Popups Only' or, if you're worried about missing the offer of a lifetime on Viagra or Webcams, set them to 'Open Popup Windows in Background'.

Mozilla Firebird

BOOKMARKS While doing without the 'Add Bookmark Here' option, Mozilla Firebird does have one ace up its sleeve in the linking stakes. The management applet is largely graphical and very easy to use; it does everything Konqueror and Opera does, but with a little more style. The killer though is the ability to open a folder full of bookmarks as tab in the same window. So, for example, your collection of bookmarks on The Roswell Incident could be opened at once for easy cross-referencing. Simply left-click on a Bookmarks folder and select 'Open in Tabs'.

HOMEPAGE Mozilla Firebird is unique among the browsers we've looked at in offering the opportunity to configure multiple home pages. The preferences are available under 'Tools>Options', and you can configure Homepages in the 'General' section. By default Firebird usually (ahem) fires up with www.mozilla.org in the window. You can type your own choice to the URL bar in this

dialogue or, if you so wish, add more by separating each address with a pipe symbol (|) which, on most keyboards, is the key to the left of **Z** with the **Shift** key.

You can have as many home pages as you like, but even at high resolutions, anything beyond six is pushing it a bit.

SEARCHING Firebird has a hybrid of Konqueror's keywords and Opera's embedded search toolbar, however the latter can also search for a text string on the currently displayed page. Type your search string and hit **Enter**.

SPOOFING Firebird relies on an Extension to change the User Agent string. To get the required extension, select 'Tools>Options' and then hit the 'Extensions' tab. Installed extensions are listed in the top pane and clicking on one will show a preview in the lower one. If you've just installed Firebird, you'll see that there are no extensions listed, so click on the blue 'Get Extensions' link, which will open a new window at <http://texturizer.net/firebird/extensions.html>. You'll find two User Agent configuration tools here, and the most useful is the User Agent Toolbar Widget by Chris Pederick. To install it, click the 'Click Here' link and then restart your browser.

Finally, do 'View>Toolbars>Customize', then drag-and-drop the UA String Widget anywhere onto the bookmarks toolbar. Hit 'Done' and you're ready to change the browser ID on-the-fly. Unfortunately the underlying identification isn't as robust as Konqueror's so, again, your success at spoofing particular webpages may vary.

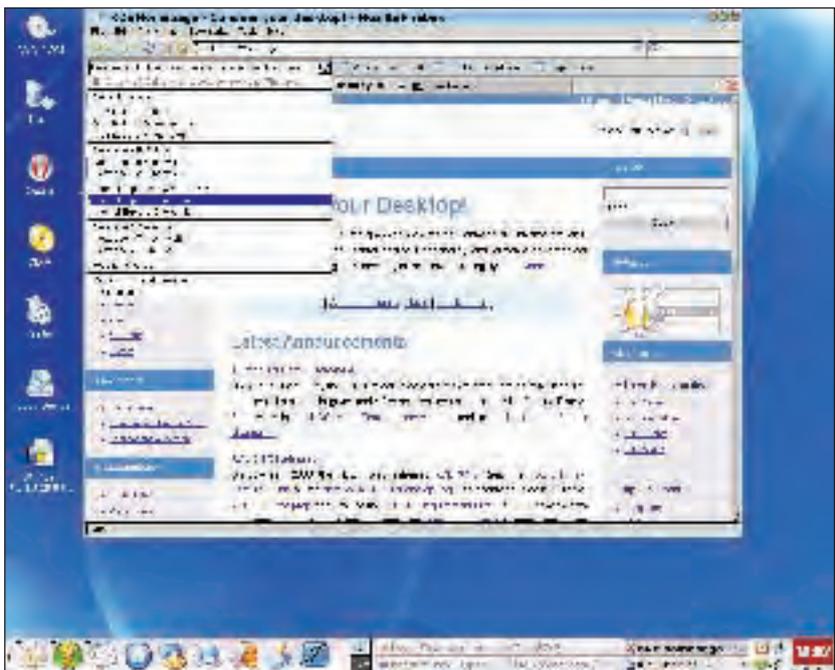
POPUIS Firebird has excellent popup blocking tools, especially the ability to allow them from specific domains – useful if your favourite sites launch popups for authentication or log in. You access these tools via the 'Tools > Options' dialog, under the Web Features tab. Selecting 'Block Popup Windows' will do just that for every site you visit, unless it is listed in the box below. You can add a site to 'unblock' by selecting 'Add site...' and typing the URL.

In this dialog you'll also find a pretty good (though sweeping) ad-blocker, which can be set to only show images from the originating website only, and as most banner ads come from a central server, they are not displayed. A great feature! **LXF**

NEXT MONTH

Next time, we move behind the firewall (you have got a firewall haven't you?) – to look at internal networking. We'll cover file sharing, Internet connection sharing and attempt to uncover the mysteries of Samba without drowning in gobbledegook.

Later in the series we'll return to the subject of browsers and, more specifically, the plugins which allow you to access 'rich media' on the web.



Unfortunately Mozilla's User Agent tool is a bit hit-and-miss on IE-centric websites.

CODE PATTERNS AND SEQUENCES

Regular Expressions

What are they, and why are they so important, you may well ask.

Dr Chris Brown has all the answers you may wish for, and more besides...

Why did you buy this magazine? Go on, be honest. It wasn't this tutorial about regular expressions, was it? No, I thought not. Regular expressions are syntactically dense, quite hard to learn, and some may think, downright cerebral. Not the kind of thing, in fact, to persuade you to part with your cash in a newsagent.

And yet, it would be hard to overestimate the importance of regular expressions in the UNIX/Linux world, and beyond. They crop up all over the place. Classically, the UNIX program *grep* searches files and prints out lines that match a given regular expression. (According to UNIX folklore, *grep* is short for *global regular expression print*.) You can use them to search for and replace text in editors like *vi* and *Emacs*, and you can use them to do automated, fancy text substitution and manipulation using programs like *awk* and *sed*. They show up in the URL-rewriting directives in *Apache*. You can use them to help select data from a relational database such as MySQL. Regular expression libraries exist for many languages, including Java, PHP and Microsoft's .NET framework. Last but not least, the technology of regular expressions powers that doyen of text manipulation languages, Perl.

Regular Expressions Defined

Regular expressions (we'll call them 'regexes' from now on) are textual expressions that describe patterns or sequences in text strings. The text strings might be lines of text in a file, or text being read from a pipe, or text held in a variable in a programming language like Perl.

Regexes are not the same as the wildcards (also known as 'filename metacharacters') that are used by the shell to match patterns in file names. For example, in the shell command:

```
$ rm [A-Z]???
```

the sequence of wildcards matches all file names that begin with a capital letter and have a three-character extension.

Syntactically, filename wildcards bear only a superficial resemblance to regexes. And whereas wildcards are only used to match against file names, regular expressions can be used to match against any text you like.

Unfortunately, there's no universally accepted definition for exactly what notations are accepted in regexes, although there is a POSIX standard, which serves as some sort of baseline. Originally there was a distinction between 'basic' regexes (as understood by *grep*) and 'extended' regexes (as understood by *egrep* = *extended grep*). But later versions of *grep* (including the GNU version which ships with Linux) understand the 'extended'

expressions, too, and other utilities add their own notations to the set and introduce many subtleties into their behaviour. For the sake of having some sort of definition to work to, this tutorial covers the extended regular expressions described in the man page for the GNU version of *grep*.

Let's get down and dirty

The table **Fig1** on the opposite page shows most of the regular expression syntax. Don't try to understand it all at once; we'll illustrate with some examples. First we need some input to process. Here's a list showing the name, maker, date of sale, and price of the most expensive scientific instruments auctioned at Christie's. We'll call the file 'auctions'.

```
Pair of globes : Mercator : 1991 : 1,723,755
```

```
Armillary Sphere : ? : 1997 : 1,252,145
```

```
Astrolabe : Habermel : 1995 : 838,350
```

```
Astrolabe : Arsenius : 1988 : 645,645
```

```
Armillary Sphere : ? : 1998 : 443,398
```

```
Rectilinear Dial : ? : 1998 : 370,614
```

Let's see how we can use *grep* to search this file. The basic syntax of *grep* is:

```
$ grep [-E] 'regex-pattern' file
```

which will show all the lines in the file that contain a match for **regex-pattern**. The optional **-E** flag is needed if we want to include extended regular expressions.

The simplest regex patterns are just literal text, so we can find how many Astrolabes were sold with:

```
$ grep 'Astro' auctions
```

which will match lines 3 and 4. We can make the search more precise by anchoring the match to the beginning of the line, using **^** like this:

```
$ grep '^Astro' auctions
```

This anchor would prevent us from getting a false match later in the line – for instance, if one of the instrument makers was called 'Astrodamus'.

Character Classes

The use of **[...]** to define character classes is a powerful feature of regexes. For example, **[Tt]** matches an upper-case or lower-case **t**, and **[aeiou]** matches any lower-case vowel. Putting **^** at the beginning of the list of characters reverses the test, so **[^aeiou]** matches anything except a vowel. You can also put a range of characters in the list, for example **[0-5]** matches a digit between 0 and 5, **[a-z]** matches any lower case letter. You can get really fancy with this, for example **[aeiouA-Z0-5]** matches a lower-case vowel or any uppercase letter or a digit between 0 and 5. Understand, though, that however much stuff is inside the square brackets, it only matches one character in the target text.

Using character classes, we can search our auctions file for sales between the years 1991 and 1996 using:

```
$ grep '199[1-6]' auctions
```

A note about notation

To help readers identify them when they occur, where I've written regular expressions with the text of this article, they are shown in bold and

coloured red. The figures **Fig1** to **Fig4** that are referred to in this tutorial are also in the Magazine section of this month's *LXF* coverdisc.

NOTATION	DESCRIPTION	EXAMPLE	WHAT THE EXAMPLE MATCHES
SINGLE CHARACTER MATCHES			
a	'Ordinary' characters match themselves	apple	The string 'apple'
[...]	Any character enclosed in []	[02468]	Any even digit
[^...]	Any character not in set	[^13579]	Anything except an odd digit
[x-x]	Range of characters	[A-Z]	Any upper-case letter
.	Any single character	c.t	cut, cat, c9t, and so on
ANCHORS			
^	Beginning of line (start of string)	^[0-9]	Lines starting with a digit
\$	End of line (end of string)	/bin/sh\$	Lines ending with /bin/sh
REPETITION MODIFIERS			
*	Zero or more of the preceding item	[a-z]*	Any sequence of lower-case letters, including none at all
?	Zero or one of the preceding item	https?://	http:// and https://
+	One or more of the preceding item	T+	T, TT, TTT, TTTT and so on
{n}	Exactly n occurrences of the preceding item	[0-9]{3}	A sequence of exactly three digits e.g. 123, 644, 999
{n,}	n or more occurrences of the preceding item.	0{3,}	000, 0000, 0000 and so on
{n,m}	Matches between n and m occurrences of the preceding item	[A-Z]{2,3}	Things like AB, ABC, YY, ZZZ
BACK REFERENCES			
\n	Matches the substring previously matched by the nth parenthesized subexpression of the regex.	((0-9))0\1	Things like 000, 101, 202, 303, 404 and so on
NAMED CHARACTER CLASSES (not the complete list)			
[:lower:]	Lower case letters		
[:upper:]	Upper case letters		
[:digit:]	A decimal digit; same as [0-9]		
[:xdigit:]	A hexadecimal digit	[:xdigit:]{4}	A four-digit hexadecimal number
[:punct:]	A punctuation character	[:punct:]	Any of: ! " # \$ % & ' () * + , - . / : ; < = > ? @ [\] ^ _ ` { } ~
[:space:]	A whitespace character	[:space:]	Tab, newline, vertical tab, form feed, carriage return, and space

Fig1 Regular Expression syntax – a quick summary.

but don't be tempted to think that you could, for example, search for dates between 1996 and 2003 using something like **[1996-2003]**. It just doesn't work like that – the regular expression engine just treats text as a series of characters, it doesn't interpret them as decimal numbers.

Regular expressions can be concatenated (placed one after the other). You've probably taken this for granted, because it's implicit in the examples we've already seen. The resulting regular expression, to quote *grep*'s man page, "matches any string formed by concatenating two substrings that respectively match the concatenated subexpressions". Got that? Well, it's obvious really – **cat** matches a **c** followed by an **a** followed by a **t**, and (perhaps less obviously) **[a-z][0-9]** matches a **c** followed by a lower-case letter followed by a digit.

Repetition Modifiers

Repetition modifiers provide a way to specify that a regular expression must match a specified number of times. This is where it starts to get interesting. The modifiers listed in **Fig1** are hopefully fairly self-explanatory. The most important thing to realise is they don't make sense by themselves, but only in combination with the regex which they follow. **+** is very useful; for example **[0-9]+** means 'a sequence of one or more digits'. The notation ***** (full point followed by an asterisk) is a common idiom which deserves special mention. **.** matches any single character, and ***** means 'zero or more of whatever the preceding item matches', so **.*** means 'any amount of anything'. Notice that ***** doesn't mean anything by itself. It's a repetition modifier and only makes sense in combination with whatever precedes it. This is quite different from the meaning of ***** in filename wildcards.

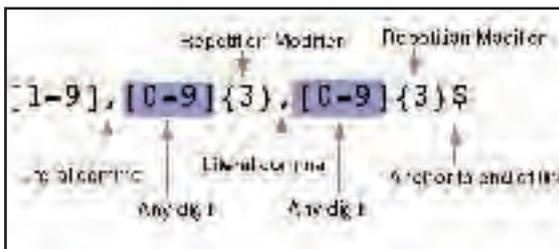


Fig2 Regexp for 'Were any items sold for more than \$1m?'

As an example, how many astrolabes were sold in 1995? For this we need a regex that matches lines beginning with **Astro** and with the string **1995** further along in the line. This following line will do the job:

```
$ grep '^Astro.*1995' auctions
```

Next is a more ambitious example. Were any items sold for more than 1 million dollars? We've already said that we can't do arithmetic with regexes, but it is possible to devise a regexp which will match a pattern like 1,234,567 but won't match 234,567. The pattern might be specified like this: a digit between 1 and 9, a comma, three digits, another comma, three more digits. We might also use a **\$** anchor to insist that the pattern comes at the end of the line. In regex notation, this leads us to the solution:

```
$ grep -E '[1-9],[0-9]{3},[0-9]{3}$' auctions
```

which is also annotated with explanations in **Fig2** above.

Grouping text

We can group parts of a regular expression together by putting them inside parentheses, in rather a similar fashion to how we can use parentheses to alter the order of evaluation in an arithmetic expression like $(3+4)*5$.



TUTORIAL Regexp

For example, **(cat)+** matches the strings `cat`, `catcat`, `catcatcat` and so on, whereas **cat+** matches `cat`, `catt`, `cattt` etc. Returning to the previous example of matching numbers like `1,234,567`, we can use grouping to write the pattern more compactly by observing that the pattern **[0-9]{3}** appears twice. By putting this pattern into parentheses, we can apply a **{2}** repetition modifier to the pattern as a whole, like this:

```
$ grep -E '[1-9]([0-9]{3}){2}$' auctions
```

Eagle-eyed readers might complain that these regexes wouldn't match a string like `20,345,678` (because we're insisting on a non-zero digit before the first comma) and of course you'd be right – it would be possible to design a 'better' regex. In fact, there's almost always a trade-off between writing a fairly simple regex that works most of the time, and a more complex one that correctly handles special cases. You're probably already beginning to realise that the syntax can get very dense very fast.

Grouping part of a regex with parentheses also gives us a way of 'capturing' a piece of text (the text that matches the grouped regex) for later use. Plucking out text patterns in this way, and subsequently storing, re-using, re-ordering or re-formatting them is at the heart of regex-driven programming. We'll meet a simple example shortly when we discuss text substitutions using `sed`, but right now let's see a simple way of exploiting grouped regexes, using back references.

Back references in regexes allow us to say "something which has been matched by a previous regular expression must appear again in the target text". The regex you want to make a back reference to must be grouped in parentheses. The notation **\1** refers back to the first grouped regex, **\2** refers back to the second one, and so on. For example, **([a-z])\1** means "any lower case letter, followed by the same lower case letter. So the command:

```
$ grep -E '([a-z])\1' auctions
```

shows all lines in the file with a repeated letter.

Quick Quiz

As you'll see, there's a certain intellectual pleasure to be gained from regular expressions, which appeals to many puzzle-solvers, crosswords addicts, and programmers. If you have access to Linux and a word list file, try the following:

1. Are there any 7-letter palindromes?
2. Are there any words containing a 'q' NOT followed by a 'u'? What about words that end in a 'q'?
3. Suppose you had a configuration file that consisted mostly of comments (lines beginning with a '#' character). Can you devise a regular expression that filters out the comment lines?

Classic Examples

With the basic syntax of regular expressions under our belt, let's enjoy three classic examples of regexes at work. All of them assume the presence of a word list, which is `/usr/share/dict/words` on my Red Hat installation. This file contains a long list of English words, one per line. In these examples we refer to the file simply as 'words'.

When pocket calculators came onto the scene, it was briefly fashionable to turn them upside down and see how many words you could spell out on the seven-segment numeric display – words like "SHELL" and "BIBLE". The geeks amongst us turned to regular expressions for an answer. First, you need to figure out which letters can be formed on an upside-down seven segment display. The answer (with a bit of imagination) is B, E, H, I, L, O, and S. So we need a regex which matches lines comprising any sequence of these characters. As a first attempt we might try:

```
$ grep '[behilos]*' words
```

but that doesn't quite work, because it admits words that contain letters not in the `[behilos]` set before or after the match. In fact EVERY line matches 'zero or more characters in the set `[behilos]`'. Anchoring the match to both the beginning and end of the line fixes the problem, like this:

```
$ grep '^([behilos])$' words
```

Now suppose we only want longish words, say seven or more letters. Instead of using the repetition modifier `*` (which means 'zero or more') we can use **{7,}** which means 'seven or more'.

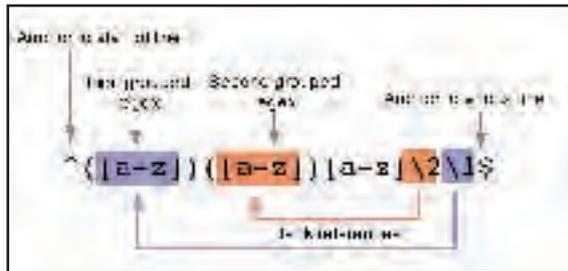


Fig3 Backreferencing annotated and explained.

Hence the following code:

```
$ grep -E '^([behilos]{7})$' words
```

which yields just five words in total: `bellies`, `hellish`, `hobbies`, `hobbles` and `lobbies`.

Our second example answers an interesting (though not terribly useful) question: are there any words in English that have all the vowels in alphabetical order? A first attempt at a regex to search the word list for such words might be:

```
$ grep '^a.*e.*i.*o.*u.*' words
```

which yields three words: `adventitious`, `facetious`, and `sacrilegious`. Close, but no cigar. The problem is that `.*` means 'any amount of anything' and we really need to say 'any amount of anything except a vowel' which is represented by the expression **[^aeiou]***. The correct solution is straightforward, if maybe a little clumsy:

```
$ grep '^([aeiou]*[a[^aeiou]*e[^aeiou]*i[^aeiou]*o[^aeiou]*u[^aeiou]*)$' words
```

which correctly returns the one word 'facetious'. (OK, don't write in, I know... there is actually another candidate, 'abstemious', but it's not in my word list file that I used for this tutorial.)

Our third challenge is to search for palindromes (words which are the same forwards and backwards, such as 'rotor'). Stated in words, the pattern we need (for five-letter palindromes) is: "any letter, any letter, any letter, whatever the second letter was, whatever the first letter was". For this we need to use backreferences, like this (also see Fig3 above):

```
$ grep -E '^([a-z])([a-z])[a-z]\2\1$' words
```

Backreferences dramatically increase the number of patterns which the regular expression engine has to test, and they can be slow. This example took two minutes to run, and yielded the list: `civic`, `level`, `madam`, `radar`, `refer`, `rotor`, `sexes` and `solos`.

Using regexes to Manipulate Text

Our examples using `grep` that have been used so far in this tutorial served to introduce regex syntax in a simple context, but they don't fully show the power of the technology. Pretty much all that `grep` can do with regexes is print out the matching lines. Using regexes within a programming language we can search text, pluck patterns out of strings, re-order them, substitute other text for them, if fact do whatever we want.

Within this short tutorial, we're not going to write any Perl code, or show Java's regex library in action, but we will at least show how regexes can be used to automate text substitutions using the stream editor, `sed`.

A quick intro to sed

The program `sed` is a stream editor. It reads a specified file line by line (or reads its standard input if no filename is given) and performs a set of editing operations on the line as specified by a `sed` 'program'. The resulting lines are then written to standard

output – they are not written back to the original file. We could easily fill an entire tutorial describing *sed*, but for now I just want to focus on one command – the substitute command. The simplest invocation of this is a command such as:

```
$ sed 's/oldpat/newpat/' file
```

which reads the named file, substitutes **newpat** for **oldpat** on every line in the file, and writes out the result. For example, in our list of auctioned scientific instruments we used **?** to indicate that the maker of an item was unknown. We could change each of these to **unknown** with the command:

```
$ sed 's/?/unknown/' auction
```

Again, I want to stress that this doesn't actually change the file **auction**. If we had wanted to do so, we would need to do something like this:

```
$ sed 's/?/unknown/' auction > temp
```

```
$ mv temp auction
```

Enough *sed*.

Now, this gets more interesting when you discover that the 'old pattern' can be a regex. Consider the deceptively simple command below:

```
$ sed 's/:.*//' auctions
```

Here, the 'old pattern' is **:.*** and the 'new pattern' is empty. So whatever the **:.*** matches is removed from the line. So the question is, just how much does **:.*** match?

Greedy Regular Expressions

In our *grep* examples, the question "Exactly how much does the regex match?" didn't matter. We were only interested in whether there was a match or not. Now it matters. The answer is that regexes are 'greedy' – they start matching as soon as they can and they go on matching as long as they can. Another slightly more precise way of saying it is that regexes return the 'left-most longest' match. To explain this in detail, consider matching the regex **:.*** against the string 'abc:def:ghi'. The regex is interpreted as 'a colon, followed by any amount of anything'. The left-most match of this is just the first colon in the string. A longer match is 'd' and an even longer one is 'de', then 'def' and so on. The left-most longest match is the string 'def:ghi' – ie from the first colon to the end of the line. So now we can see that the previous 'deceptively simple' example

```
$ sed 's/:.*//' auctions
```

strips out everything except the first field of each line, leaving only the instrument names.

Here's another example, subtly different:

```
$ sed 's/:.*//' auctions
```

Here, the regex **:.*** means 'any amount of anything followed by a colon'. What is the left-most longest match for this in the lines in the auctions file? Everything up to and including the last colon, of course. So the command strips out everything except the last field, leaving only the prices.

Back references in substitutions

Let's take our use of regexes in *sed*'s substitute command one level further. The challenge is to replace occurrences of strings like **\$50** with **50 dollars** so that this input:

```
$300 and up is considered expensive,
```

```
but most people will happily pay $50 to $100
```

comes out as:

```
300 dollars and up is considered expensive,
```

```
but most people will happily pay 50 dollars to 100 dollars
```

This isn't a simple replacement of a **\$** by the string **dollars**. We

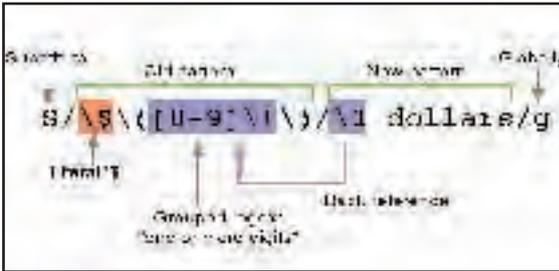


Fig4 Prefacing characters with a backslash.

must find occurrences of a **\$** sign followed by a string of digits, and replace it by the same string of digits followed by the word **dollar**. In *sed*, as in *grep*, we can use parentheses to group parts of the regex, and the backreference notations **\1**, **\2** and so on to refer back to these in the 'new pattern' part of the substitution. The syntax gets quite messy, as we'll see. A simple regular expression to match a string of digits might be **[0-9]+**, so our first attempt might look like this:

```
$ sed 's/$([0-9]+)/\1 dollars/' tfile
```

but it turns out that there are a couple of syntactic nuances we need to get right before this will work. For a start, the **\$** is interpreted to mean 'end of line' which isn't what we want here. To turn off the special meaning of the **\$**, so that it represents a literal **\$** symbol, we 'escape it' by putting a backslash in front like this:

```
$ sed 's/\$([0-9]+)/\1 dollars/' tfile
```

...but this still doesn't work. The problem is that parentheses and the **+** metacharacter (along with several other metacharacters used in 'extended' regular expressions) are not treated as special unless they are prefaced by a backslash. This knowledge leads us to our third attempt:

```
$ sed 's/\$([0-9]+)/\1 dollars/' tfile
```

Yes, all the backslashes are confusing. The one before the **\$** is turning the special meaning of **\$** off, and the ones before the **(**, **+** and **)** characters are turning their special meaning on. **Fig4** may help to make the expression clear.

One last change

There is one last change we need which is a result of how *sed* works rather than the behaviour of the regular expression. In a substitute command, if the 'old pattern' appears more than once on the line, only the first occurrence of the pattern will be substituted. Appending a **g** to the end of the substitute command makes the substitution 'global' – ie it is performed for every occurrence of the 'old pattern' in each line. This leads us to our final solution:

```
$ sed 's/\$([0-9]+)/\1 dollars/g' tfile
```

I should stress that the **/g** notation is a feature of *sed*'s substitute command, *not* part of the regex syntax. By the way, these substitute commands also work in *vi*, where you can apply them to the current line or a specified range of lines. **LXF**

Reading

It's hard to do regular expressions justice in this short tutorial. If you want to study them further, two books are worth a mention, both from O'Reilly. First, there's *UNIX Power Tools* by Powers, Peek, O'Reilly, and Loukides ISBN 0-596-0330-7. This is an intensely practical book about using UNIX at the command line and covers many tools that understand regular expressions. Second, there's *Mastering Regular Expressions* by Jeffrey Friedl, ISBN 0-596-0289-0 – a real tour-de-force of a book which not only covers the mechanics of regexes in extraordinary detail but also talks about efficiency, and the use of regexes in Perl, Java, and .NET. It also has a huge number of examples. If you use regular expressions as part of your professional work (even if you already have a good book on whatever language you're programming in) I would strongly recommend this book to you. Happy regexing!

HOMEWORK

Find your own solution

If you'd like a final example to get your teeth into, try this. Suppose someone had reversed the order of the two arguments in the standard C function **listen**, so you needed to edit all your C source code files so that a line like

```
rv = listen(socket, 5);
```

(where **socket** and **5** are just two parameters

being passed to the function) becomes

```
rv = listen(5, socket);
```

Can you devise a substitute command in *sed* that will automate this?



FLEX & BISON

Compiler Writing

Paul Hudson keeps the pressure on by extending the compiler to handle conditional statements...

There's so much we need to cover this month that I barely have space to welcome you and discuss the story so far. So, here's a condensed version: our compiler, SKYlang, reads, compiles, and executes scripts with variables, basic operations, and loops. It also handles intermediate code and type conversion, although there are a few (known) bugs in the system right now.

We're long overdue a "catch-up" issue where we can track down bugs in the code and fix them properly, but right now I think you're all itching for new features – I certainly am! So, this issue we're going to build upon the looping code we looked at last issue, and use the same method to make conditional statements work.

Basic code changes

Add this code after the **T_ECHO** match in skylang.l:

```
"if" {
    return T_IF;
}
```

Now add this code after the **T_DIVIDE** match in the same file:

```
"%" {
    return T_MODULO;
}
```

```
"==" {
    return T_ISEQUALS;
}
```

Finally, add this code after the **T_CBACE** match, also in skylang.l:

```
"(" {
    return T_OBRACK;
```

```

}
)" {
    return T_CBACE;
}

```

As you can see, we're going to make our compiler understand modulo and testing whether two variables are equal. Add the operations **SOP_JMPEQ**, **SOP_JMP** and **SOP_MOD** to the `SkyOpType` list in `skylang.h` – this allows us to jump if equal, jump unconditionally, and perform modulo arithmetic. You'll need to add **T_IF**, **T_OBRACK** and **T_CBACE** to the `%token <var>` list in `skylang.y`, and also add **T_MODULO** to the `%left` list a few lines below.

Add this code after the **T_DIVIDE** `sky_do_binop()` code:

```
| expression T_MODULO expression {
    $$ = sky_do_binop(SOP_MOD, $1, $3);
}

```

The **SOP_MOD** code for our `execute()` function is basically the same as that for **SOP_ADD**, with the exception that modulo arithmetic only works on integers, which means we need to convert floats to ints where necessary. Here's how that looks:

```
case SOP_MOD:
    if (op->op1->type == svString) { op->op1 = sky_strtoint(op->op1); }
    if (op->op2->type == svString) { op->op2 = sky_strtoint(op->op2); }

    if (op->op1->type == svInt && op->op2->type == svInt) {
        double dtemp = op->op1->intval % op->op2->intval;

        if (dtemp > INT_MAX || dtemp < INT_MIN) {
            op->result->type = svFloat;
            op->result->floatval = dtemp;
        } else {
            op->result->type = svInt;
            op->result->intval = op->op1->intval % op->op2->intval;
        }
    } else if (op->op1->type == svFloat && op->op2->type == svFloat) {
```

```

op->result->type = svFloat;
op->result->floatval = (int)op->op1->floatval % (int)op-
>op2->floatval;
} else if (op->op1->type == svFloat && op->op2->type ==
svInt) {
op->result->type = svFloat;
op->result->floatval = (int)op->op1->floatval % op->op2-
>intval;
} else if (op->op1->type == svInt && op->op2->type ==
svFloat) {
op->result->type = svFloat;
op->result->floatval = op->op1->intval % (int) op->op2-
>floatval;
}
break;

```

Note the **(int)** typecasting scattered throughout – there are faster ways of converting to integer, but they are invariably more complicated to implement.

Conditional statements

With the basic changes out of the way, we can concentrate on the focus of this tutorial: adding conditional statements. So far, we haven't actually provided a definition of what a SKYlang **if** statement looks like – so here it is:

```

if ($foo == $bar) {
SKYlang code here
}

```

Naturally it should be possible to have other operators in there other than just **==**, as well as other operands. Indeed, it should be possible to have any expression in place of the operands, such as **if (1 + 1 + 1 == \$foo)**. It should also be possible to have an optional else statement, like this:

```

if ($foo == $bar) {
SKYlang code here
} else {
SKYlang code here
}

```

However, adding **else** is something we'll leave for later on – for now, we're going to concentrate on basic conditional statements using the **==** operator, "is equal to".

As with loops, there are several golden rules to if statements that we will consider absolutely inviolate:

- i If statements must have an opening and closing brace
- ii The closing brace always matches the most recently opened brace
- iii You can stack if statements together
- iv Any number of if statements can be stacked

Now consider how an **if** statement should be handled, remembering that we're working without the else statement right now to keep things simple. While there is still some flexibility as to how you implement this, the situation is still programmatically fairly similar to the looping system we had in *LXF45* – here's how I'd do it:

- i Check variables
- ii If they match, jump to start of code execution
- iii If they don't match, jump to end of code execution

In intermediate code, using our new operations, that would look something like this:

```

SOP_JMPEQ op1:$foo, op2:$bar, result:<op of this statement>
SOP_JMP result:<op of last statement inside brace>

```

```
<statements>
```

As with the looping in the last issue of *LXF*, we're using the result SKYVar of our operations to hold the operation number to jump to – and, just like last issue, we need to jump to the operation *before* where we want to be, because our **execute()** loop will advance the op counter by one after we set it, effectively making it point to the correct place.

Now, to match our if statement, we'll need this *Bison* code:

```

T_IF T_OBRACK expression T_ISEQUALS expression
T_CBRACK T_OBRACE statements T_CBRACE

```

As with the looping code, we want to get control immediately after the **T_OBRACE**, pass it back to parse *statements*, then get it back again after the **T_CBRACE**. The solution is the same as last time – we need to use embedded actions. Here's how the rule should look:

```

T_IF T_OBRACK expression T_ISEQUALS expression
T_CBRACK T_OBRACE { do stuff } statements T_CBRACE { do
other stuff }

```

If you recall from last issue, embedded actions are considered independent of each other – that is, we can't access the code performed in **do stuff** from **do other stuff** because they end up being different blocks of C code. Because *statements* can be an arbitrary number of lines of code, we need to store the location of our **SOP_JMPEQ** statement inside the **do stuff** block, add the *statements* to our opcode list, then retrieve the **SOP_JMPEQ** statement and set its return value to the current position in the **do other stuff** block – this allows us to work around the embedded actions gap, and also to stack up if statements.

The solution is the same as we had for loops – we use an STL stack of integers that hold the operation numbers to jump to, pushing and popping as necessary. We also need to make sure that we jump if the conditional statement returns false, so this is a little more complicated than loops in some respects.

Here we go, then

Just below the definition of **LoopStack** in *skylang.y*, add this line:

```
std::stack<int> IfStack;
```

This is where we'll be storing the jump statements for our stack, just as with our loop stack. Now, just after the **ECHO** action in *skylang.y*, add the following code: (NOTE: I've included the last couple of lines from **ECHO** to help you match it up)

```

new_op->op1 = $2;
}
|
T_IF T_OBRACK expression T_ISEQUALS expression
T_CBRACK T_OBRACE {
SKYOp *new_op = new SKYOp;
OpArray.opcodes.push_back(new_op);

new_op->opcode = SOP_JMPEQ;
new_op->op1 = $3;
new_op->op2 = $5;
new_op->result = new SKYVar;
new_op->result->intval = OpArray.opcodes.size();

SKYOp *uncond_jump = new SKYOp;
OpArray.opcodes.push_back(uncond_jump);

uncond_jump->opcode = SOP_JMP;
uncond_jump->result = new SKYVar;

```



TUTORIAL Compiler

```

<< IfStack.push(OpArray.opcodes.size() - 1);
} statements T_CBRACE {
int opnum = IfStack.top();
IfStack.pop();
OpArray.opcodes[opnum]->result->intval =
OpArray.opcodes.size() - 1;
}
|

```

Although the individual lines are highly similar to the looping we covered last issue, the structure is quite different because of the kind of intermediate code we want to generate. As mentioned earlier, we need two jumps to handle a conditional statement properly. Consider the following script:

```

$foo = 1;
$bar = 2;
if ($foo = $bar) {
echo $foo;
}
echo $bar;

```

Here's roughly how that looks in mock intermediate code:

```

1 ASSIGN $foo = 1
2 ASSIGN $bar = 2
3 JMPEQ $foo, $bar, 5
4 JMP 6
5 ECHO $foo;
6 ECHO $bar;

```

The last numbers in **JMPEQ** and **JMP** are the line numbers to jump to – as you can see, the unconditional jump skips over the contents of our braces, effectively executing none of the code. Furthermore, the **JMPEQ** statement jumps over the unconditional jump and into the contents of the braces, effectively executing the code. This is precisely how an if statement should work, and it's how I've implemented it in the source code.

The first Bison action creates two new operations – a **JMPEQ** and a **JMP**, as planned. Note that the **JMPEQ** has its result set to **OpArray.opcodes.size()** – this is key! If you recall, our opcodes STL vector is 0-based, which means the **OpArray.opcodes.size() - 1** would relate to the **JMPEQ** statement we just added. By dropping the -1, we jump to the next statement after the **JMPEQ**, which is always going to be the unconditional **JMP**, and, thanks to the fact that our **execute()** function always increments the opcode pointer one beyond where we set it with each iteration, jumping to **OpArray.opcodes.size()** will end up with us executing code from line 5 onwards of the mock intermediate code above.

We then add another operation for the unconditional **JMP** – note that we don't bother creating an **op1** or **op2** variable for it, because **JMP** *always* jumps, without needing to check any variables. Furthermore, because there can be any number of statements between our unconditional jump and the location it should jump to *and* because we don't yet know where it should jump to, we push the current operation onto our **IfStack** stack.

In the second part of our action, we now know how many statements there are inside the braces, so we can go back and fill in where the unconditional **JMP** should jump to. This is done by first taking off the top item from our stack, popping it to delete it, then setting the result of the appropriate operation to the current position.

Easy? Well, if you found last issue a doddle, this will have been no different because it builds on the same skills. However, the hard part is yet to come...

Making it execute

All we've got so far is a parser rule to handle our if statements and put them together into intermediate code – there's nothing to make them execute yet. Handling unconditional jumps is very easy because there's no checking involved, so the code should jump regardless.

Therefore, the case statement for unconditional jumps is simply this:

```

case SOP_JMP:
i = op->result->intval;
break;

```

The code for **SOP_JMPEQ** is more difficult, though – even more difficult than the code we currently have for **SOP_JMPLTE**. The reason for this is because **SOP_JMPLTE** was designed to handle very basic **for** loops – just take a look at the check it does:

```

if (UserVars[op->op1->charval]->intval <= op->op2->intval) {

```

It automatically assumes that **op1** will be a variable and **op2** will be a number, which means that checks such as **\$i <= 5** will work fine, but **\$i <= \$j** and **4 <= 5** won't work at all. This is, of course, wrong – while it's suitable for teaching how to make loops, it's not good in the long run. Because our if statement will need to be able to compare more complex values, the code to handle it needs to be much more complicated – it needs to handle strings, ints, floats, and variables interchangeably.

First things first, though – you need to edit the start of your **execute()** function to look like this:

```

void SKYOpArray::execute() {
int abs_op_count = 0;
SKYOp *op;

SKYVar *tmpop1;
SKYVar *tmpop2;
bool jump;

for (int i = 0; i < OpArray.opcodes.size(); ++i) {
++abs_op_count;
jump = false;
op = OpArray.opcodes[i];

if ((op->op1 != NULL) && op->op1->type == svString) {
if (UserVars[op->op1->charval] != NULL) {
tmpop1 = UserVars[op->op1->charval];
} else {
tmpop1 = op->op1;
}
} else {
tmpop1 = op->op1;
}

if ((op->op2 != NULL) && (op->op2->type == svString)) {
if (UserVars[op->op2->charval] != NULL) {
tmpop2 = UserVars[op->op2->charval];
} else {
tmpop2 = op->op2;
}
} else {
tmpop2 = op->op2;
}

switch(op->opcode) {

```

Debugging your code

If you use a debugger like *gdb* or *gvd*, you can add debugging information to your compiler code to try to figure out where things are wrong if you have a bug. To get debugging information into your executable, edit your `makefile`, and change the `CC` line to `CC = g++ -g`

New to this code are **tmpop1** and **tmpop2**, which will be temporary holding areas for converted variables for each operation. Once we get around to cleaning the code up next issue, these temporary variables will be used across the board in place of **op->op1** and **op->op2**. There's also code in there to check whether each operand in a given operation isn't **NULL**, and, if it's a variable, to set it to the value of the variable rather than using the actual variable itself. This is merely a modification of the code we originally placed into **sky_strtoint()** – doing it this way is preferred, as will be seen when we clean up the code. Note the boolean variable **jump** I slipped in there – this is used later on.

That's still not the hard bit, though – actually handling the **SOP_JMPEQ** code is the trickiest part, simply because it's more convoluted than the other pieces. As per usual I've tried to keep it simple so that you learn concepts not code – at the end of the day, most of the code is just there to handle each possible operand type.

Here's the case statement for **SOP_JMPEQ**:

```
case SOP_JMPEQ:
switch (tmpop1->type) {
case svString:
switch(tmpop2->type) {
case svString:
jump = (strcmp(tmpop1->charval, tmpop2->charval)
== 0);
break;
case svInt:
tmpop1 = sky_strtoint(op->op1);
jump = (tmpop1->intval == tmpop2->intval);
break;
case svFloat:
tmpop1 = sky_strtoint(op->op1);
jump = (tmpop1->intval == tmpop2->floatval);
break;
}
break;
case svInt:
switch(tmpop2->type) {
case svString:
tmpop2 = sky_strtoint(op->op2);
jump = (tmpop1->intval == tmpop2->intval);
break;
case svInt:
jump = (tmpop1->intval == tmpop2->intval);
break;
case svFloat:
jump = (tmpop1->floatval == tmpop2->floatval);
break;
}
break;
case svFloat:
switch(tmpop2->type) {
case svString:
tmpop2 = sky_strtoint(op->op2);
jump = (tmpop1->floatval == tmpop2->intval);
break;
case svInt:
jump = (tmpop1->floatval == tmpop2->intval);
break;
case svFloat:
jump = (tmpop1->floatval == tmpop2->floatval);
```

```
break;
}
break;
}
if (jump) {
i = op->result->intval;
} else {
// loop finished, do nothing
}
break;
```

If you hate copying so much code in at once, relax – the latest copy of SKYlang is on this issue's LXF coverdisc!

The easiest part to understand is the last part – if **jump** has been set to true, we change the main loop variable to the jump point, effectively changing the next operation to be executed. The rest of the code handles the fact that a given check might compare a string against a string, a string against a float, a float against an int, etc.

While a lot of the code is basically just duplication, there are three distinct types of check done. The simplest takes this form:

```
jump = (tmpop1->floatval == tmpop2->intval);
```

Here the **floatval** of **op1** is compared to the **intval** of **op2** – if the two match, it will set **jump** to true.

```
jump = (strcmp(tmpop1->charval, tmpop2->charval) == 0);
```

Here we use the **strcmp()** function, which compares two strings, to compare our two **charvals** – if the first string is 'less' than the second string, **strcmp()** returns **-1**, and if the first string is 'greater' than the second string, it returns **1**. Two strings that are the same will return 0, although they do need to be an exact match. Note that 'less' means 'before in the alphabet', whereas 'greater' means 'comes after in the alphabet'.

```
tmpop1 = sky_strtoint(op->op1);
```

```
jump = (tmpop1->intval == tmpop2->intval);
```

This last check compares a string to an integer. There are two ways this could have been done – we either convert the string to an integer then compare it against the integer operand, or we convert the integer to a string, then compare it against the string operand. Both are possible with little effort, but the chances are that the former will be more useful to users. So, that code block runs our **sky_strtoint()** function, then sets **jump** to true if the two operands match.

Conclusion

This has been another "features drive" issue, where we've added lots of new functionality without going back to fix what's there already. While this is good for teaching because it keeps pushing you onwards to try new things, it's not ideal for our language – it's got quite a few bugs and holes that need to be corrected!

For example, I haven't yet implemented **T_DECREMENT** to handle decrementing loops; I haven't forced loops to jump if they don't match the first time around; if you say, echo "hello", you get a segfault because **SOP_ECHO** only works with variables right now; also, wouldn't it be nice to unify the code for our comparison operators so that they all worked as **SOP_JMPEQ** does? This is all for next month, though – for now, I suggest you grab the code from the coverdisc and make sure your own copy matches the important parts. There isn't room for me to highlight all the code changes here, but why not visit the forums at www.linuxformat.co.uk and discuss them with other readers? 



NEXT MONTH

Next issue we're going to do some vacuuming of our compiler, particularly with the focus on making it work smarter by trying to cut down the amount of code repetition in there. This will make life much easier later on, and to be honest, it's long overdue!

ILLUSTRATION BY HIROSHI SAITO - www.blender3d.org

ADVANCED EFFECTS

Create a Particle system

PART 7 Stepping sway from modelling, **Jono Bacon** looks at how to use a system that can be used to create such impressive things as smoke effects.

This month we are going to take a look at an area in *Blender* that is completely new to us; particle systems. If you have not used a particle system before in a 3D modeller, you are probably wondering what I am rambling on about. Fear not, I will explain.

The *Blender* particle system is an incredibly powerful means of creating natural phenomena and effects. The typical example where particle systems are used is in creating fire. If we were to

attempt to create fire with normal shapes and textures, we would not be able to create something suitably realistic, and it would look distinctly artificial, spoiling any atmosphere that you had laboured to create in the rest of your scene.

Particle systems let us get around this limitation of traditional modelling, allowing us to create effects where a large number of small particles are spat out from a source. This system can be used for phenomena such as fire and smoke as described, but it

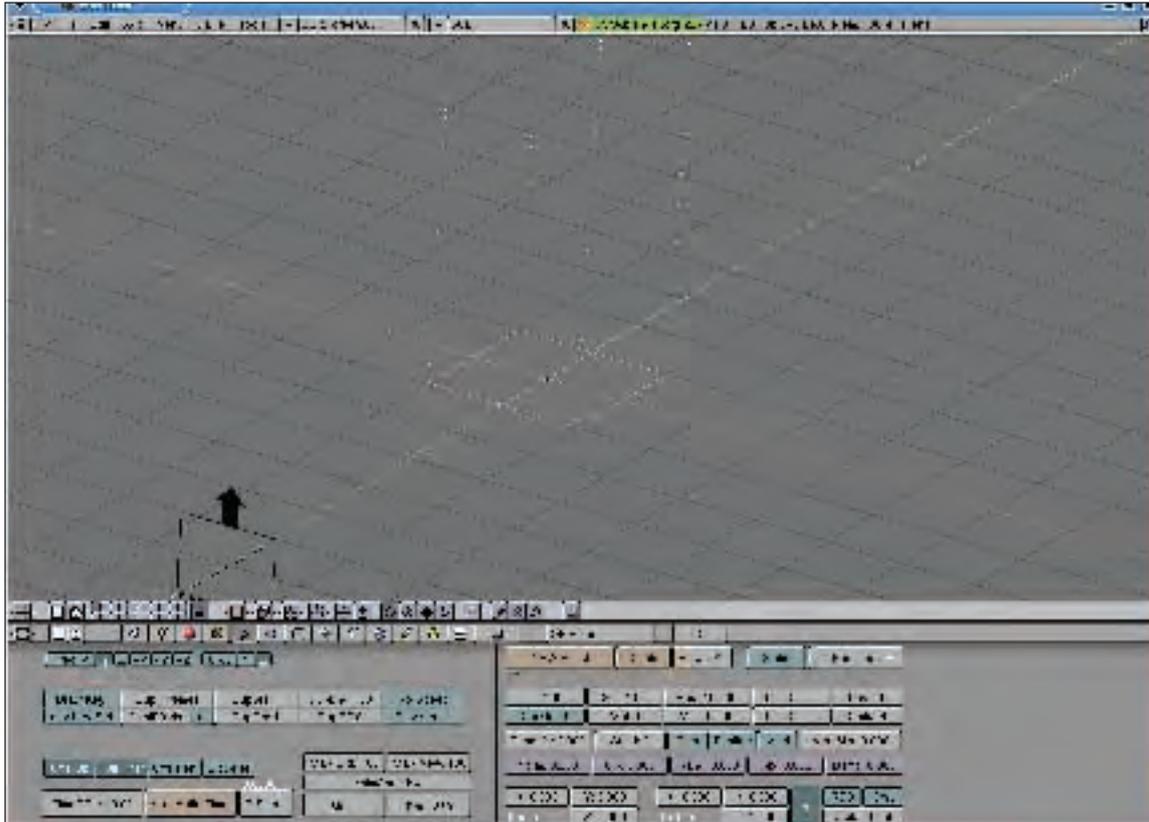


Fig1 A first particle system spitting out a number of particles from the corners of the plane.

can also be used to create large numbers of single objects such as a flock of birds or a flurry of ants.

This month we will first look at how to create a particle system, and then we will address how to create various effects and animations. Before we begin though, a word of warning. Particle effects systems require enormous computing power as each tiny particle needs to be rendered individually. It is therefore recommended that you only attempt to create particle systems that reflect your computing power. If you are on a slow machine for instance, don't create a particle system with hundreds and hundreds of particles. Experience plays a part here, so experiment with *Blender* to see exactly how much of a battering the capabilities of your machine can take.

Getting started

Particle systems have two major parts – an emitter and the particles. An emitter is similar to a fan; it pushes something out. The difference with a fan is that the air is generally blown through the fan in just one way – you cannot really control how the air is blown out; just the direction of the air by twisting the head of the fan around. With a particle system you can control the many different aspects of the emitter and add additional forces such as gravity and wind.

The whole point of a particle system is that it should be realistic in its physics if you are using it in a real world setting (such as smoke or fire), but it should also be able to have non earth-like physics if required. The flexibility is within *Blender* to modify the emitter exactly how you need it.

To get us started, open up a brand new *Blender* project. Do not remove the default plane as we normally do – this time, we will use this as the emitter. You can reasonably use any mesh object as an emitter, but bear in mind that different meshes will have different effects.

Now, using both mouse buttons together, rotate your view and use the zoom button to get a good angle on the default plane. Now click on the Animation Button (**F7**) and click on the New Effect button. You will now see a drop down box to the right with Build written on it. Click on this box and select Particle from it. You will see a bunch of buttons appear that are all used to configure your particle system.

These buttons have a variety of different settings, but the most basic button to create some particles is the Norm button. If you adjust the value of this button and press **Alt-A** (with the mouse over the 3D view) to preview the animation, you will see the particles coming from the emitter which is the default plane. You can also adjust how many particles are being generated by adjusting the Tot button.

With our system now working, we can now affect the physics of the particles that are getting emitted. The first thing you may want to do is have some wind blowing the particles. This can be configured with the Force buttons. You will see that there is a



USES FOR PARTICLE SYSTEMS

Repeating object realism

You may be wondering what exactly you can use these particle systems for. Although the obvious options are smoke and fire, there are many other effects you can develop. Anything with a lot of repeating objects can be created such as grass, armies of soldiers, pebbles, gravel etc. Also remember that you can create smog, fog, rain, hail and snow using particle systems.

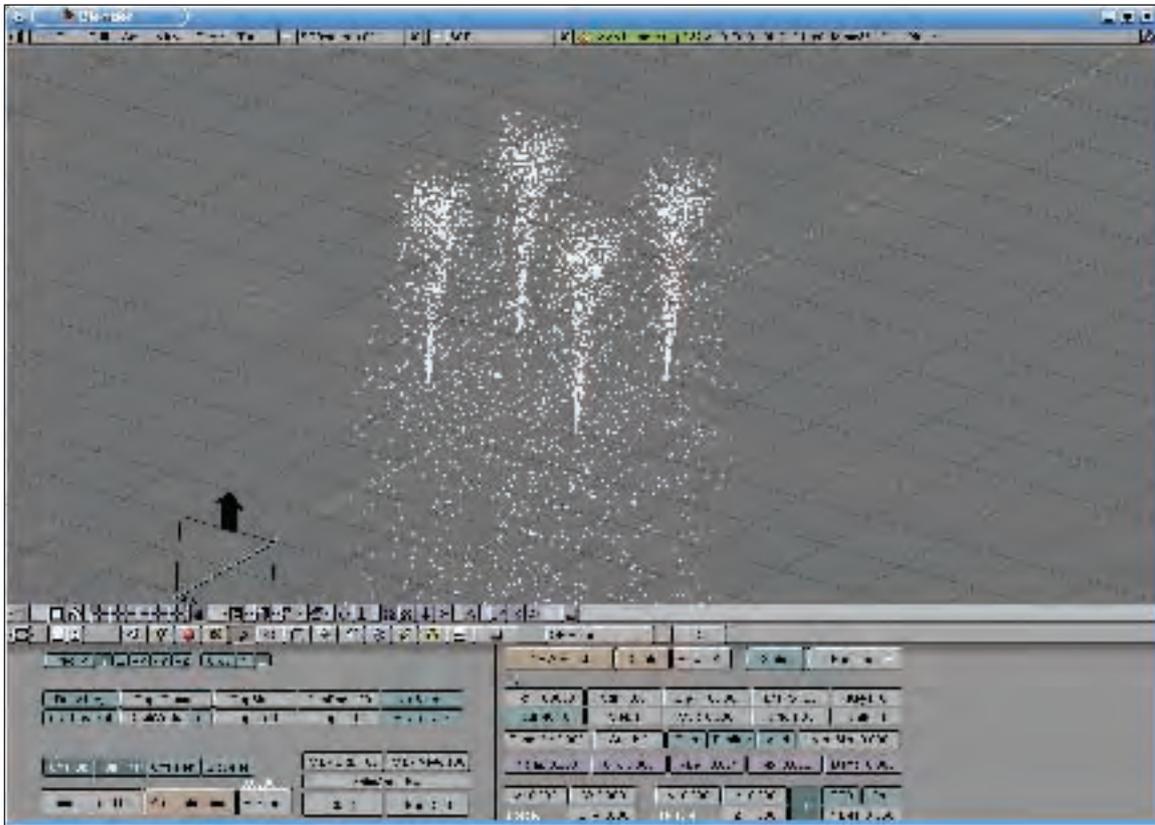
Particle systems are often used to add atmosphere to a scene. People know that 3D scenes are comprised of shapes, and when you

add something that appears to have no set shape such as fire, it makes the scene look more natural. If you have a normal scene for example, and then place a particle emitter just below the camera with some transparent fog to add some atmosphere.

As with anything to do with creative tools such as *Blender*, it is best to think of what you want to do and then figure out how to make *Blender* do it. Knowing how to create a particle system, allows you to practically implement these effects in your scenes.

TUTORIAL Blender

Fig2 Adding randomness to particles can create a showering effect.



◀ button for each axis (X, Y and Z) that can be used to affect how much wind is blown in that direction. The rule of thumb as normal is to play with the buttons to see what effects you can achieve. Remember these buttons can be combined together well. You can see this effect in **Fig1** on the previous page.

A few other buttons that can help with our system include those to adjust the randomness and life of the particles. Randomness is an important factor with particles as the effect we wish to achieve is often a random phenomena such as smoke; this randomness is controlled with the Rand button. You can also configure the randomness of the particles by adjusting the Seed and Randlife buttons. If you would like to set how long each particle is alive for, you can use the Life buttons, and this can be used to control the extent of which the emitter spits out the particles.

Even at this basic level, we can create an authentic looking particle system. Set the Force buttons for X and Y to 0, but set the Z button to -1.000. Set Tot to a large number such as 80000, Rand to 0.064 and the top Life button to 55.00. When you run the animation, you should see the particles shoot up and then shower down like in **Fig2**. At this stage, we just need to actually create a rendered animation of the sequence with some suitable materials.

Rendering the system

When you move the camera and add a spot light to the emitter, you will notice when you render the scene, that it looks like it is spitting out a number of white balls. These white balls are the default materials for a particle, and as you can see, they are quite large. We will now adjust these materials to ensure they look sleeker and more how we envisage them. I would quite like to create an effect where glittery type stardust is shot out of the emitter. We can do this by using a Halo.

First, right-click the emitter (the default plane), and add a new material. In the Material Buttons there are a number of buttons down the centre of the button area. In these buttons is the Halo button; click on it, and you will see the preview window change to a small blot of colour. We can now change the shape of the halo to one that looks more suitable; we will first use the Line shape. In the middle row of the Material Buttons, click the Lines button. You can then click the Line button under the colour block, and change the colour of the lines using the R, G, B sliders. You can now use the Lines button under the R, G, B buttons to select how many lines the halo has. Now animate the sequence by pressing the ANIM button in the Display Buttons (**F10**). You will now see the halo particles flowing out of the emitter. Although using the Lines button works well, there are other effects you can use also such as Rings and Stars. See the Halo Types box to see the differences.

Creating smoke

Although we have discussed how to create a particle system, our efforts have been on creating a system that spits out a lot of small objects. If you would like to create a more continuous entity such as smoke, we will need to tweak our particle system settings a little more.

First of all, you will need to reduce the size of the emitter using the scale key (**S**). Scale the emitter so the four corners are quite close together. We now need to change our halo so it looks like a bit of smoke. To do this we will use a normal halo, so ensure the Lines, Rings and Stars buttons are not selected. Now change the colour of the halo to a grey colour, and use the Alpha and Add buttons to give the material some transparency. The previewed material should be barely visible when you have done this.

If you now click on the HaloTex button, we can apply a texture to our halo material to give it a more realistic look. When you have

clicked the button, you need to go to the Texture Buttons (**F6**) and select the Cloud texture. If you now return to the Materials Buttons (**F5**) you need to adjust the slider on the right-hand side of the buttons to change the colour from pink to a more suitable colour. This colour is the colour the texture blends in with, so it therefore needs to be similar to the general colour of your halo.

With the material set, we now need to go and adjust how our emitter works. Currently it is showering the particles down, and we need to give the emitter a smoke machine kind of effect. To do this you must first increase the particle count (Tot slider) to at least 1000. The aim of this effect to make the large amount of particles visually mold together to form a continuous smoke effect.

Another setting you should change is the Randlife and Rand buttons. Smoke is a random entity that can be blown all over the place, so we need our emitter to represent this. Set the Randlife button as high as possible, and the Rand button to about 0.50. Again, practice does not make perfect – perfect practice makes perfect, so have a play with the settings to try different effects. The final thing we need to change is the Force buttons. Adjust X, Y and Z buttons to ensure that the smoke bellows in the fashion you want it to. You may want to experiment with some of the negative values of these buttons also to create some interesting effects. You can see this kind of effect in **Fig3**.

Using objects as particles

So far, we have been using halos as our particle materials. Although this can create some impressive and useful effects, we are somewhat limited to these effects. An impressive feature in the *Blender* particle system is the ability to use any mesh or object as a particle. With this ability you can create incredibly impressive looking scenes with hundreds or thousands of objects. We will take basic look into this system and explore how to use objects as particles, but these basic techniques form the structure for more advanced uses of the system. Once again, I must warn you that the process required for steps is large, so you may want to make yourself a cup of tea while you are waiting for sequences to render. Using complex objects as particles can have incredible processing requirements, so be aware of this issue.

To get us started, create a new scene and set the default plane as the particle system. Set a reasonably low amount of particles and create low amount of randomness. Now, create an object such as a ball, cube or other shape next to the emitter and scale the object right down so it is representative of the scale of the emitter. You may need to increase the size of the emitter to achieve the correct scale.

Now select the object, and then holding Shift, select the emitter. Press **Ctrl-P** and make the emitter a parent of the object. You now need to select just emitter, and ensure the DupliVerts option is turned on. When you now preview the scene, you will see the objects coming from the emitter. A final option to turn on is the Vert option in the particle emitter buttons. This option will make the objects spin and twist in the direction of the particles. You can see the result of of my system using cones as the object in **Fig4**.

Conclusion

This issue we have taken a good look at how to create a particle system, set materials, change the shape of the particles and create smoke effects. We have then gone on to use objects as particles. These particle effects systems are very powerful tools, and can greatly enhance the realism of a scene. Until we resume next month, I recommend you spend a good while playing with

HALO TYPES

When you are using halos as your materials, there are three major types of halo that you can use:

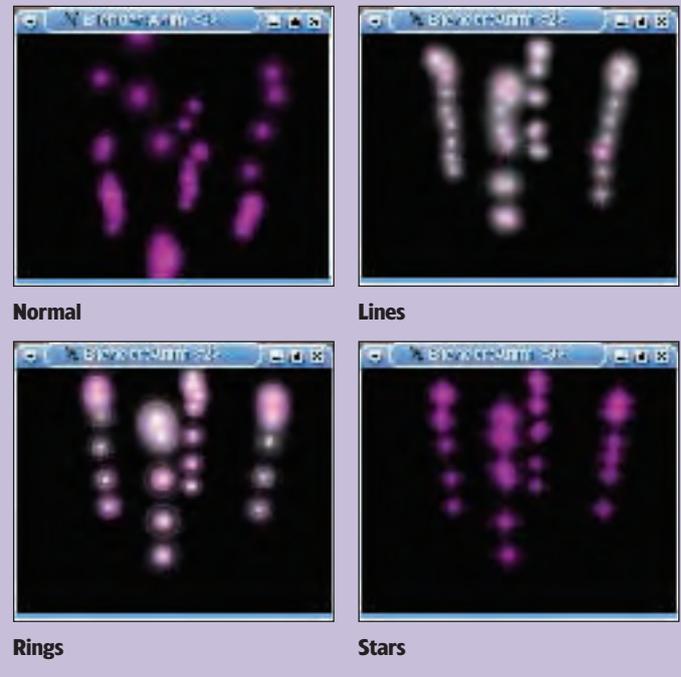


Fig3 Smoke effects can be developed by adding textures to halos.



Fig4 Adding objects as particles can create interesting looking scenes.

the different settings for the particle settings. The sheer amount of effects and displays you can achieve with these buttons is impressive, and the only way you can learn is by experimenting with them. It is also recommended that you play with adding objects as particles, and try adding different lights to create shadow effects. You may also want to add other (non-particle) objects to the scene and combine different particle systems. Finally, with all of this processing going on, you may want to go and buy a new PC or brew a whole vat of coffee... [LXF](#)

NEXT MONTH

We will continue to look at particle system effects, and learn how to create static as opposed to animated effects. In *LXF47* we will also begin looking at some more lighting effects, and touch on more techniques to add realism to your scenes.

IMAGE EDITING

GIMP photo restoration



Digital image manipulation is a subject area that reaches far beyond web banners and movie magic. *The GIMP* provides the tools necessary for quality photo restoration, as **Michael J Hammel** demonstrates.

An often overlooked corner of the digital image editing realm is photo restoration and retouching – the art of preserving and enhancing old photographs that have been folded, misused or cracked and faded with age. While the face value of most 100-year-old photographs is little more than their original worth, the personal value can seldom be measured. Restoring worn and faded photos to their original or near original look can be of high value to families and even corporations looking to track their lineage. In some cases, restoring photographs is the only way to verify facts. A damaged photo might show if an individual was present at the time the picture was taken. It could also show the time of day from the clock on a wall at an important event, or by orienting the shadows cast on the ground at the same location years later.

There are a number of complex tricks that can be used to rebuild damaged images, but the most common trick is to simply clone similar areas as replacements for the damaged areas. The GIMP provides a Clone tool that allows using brush strokes (and,

thus, any brush shape) to copy one area to another. However this method is only appropriate for small cleanup work such as removing a dust speck or hiding a thin, short scratch. Additionally, this method is destructive – the cloning occurs on the layer where the damage exists so if your patches are not to your liking you may not be able to undo them easily.

A better approach for larger blemishes is to make a selection, copy and paste it as a patch layer and blend this patch layer into the original using the airbrush tool in a layer mask. This allows changes to be made later either by modifying the layer mask or completely replacing the patch layer.

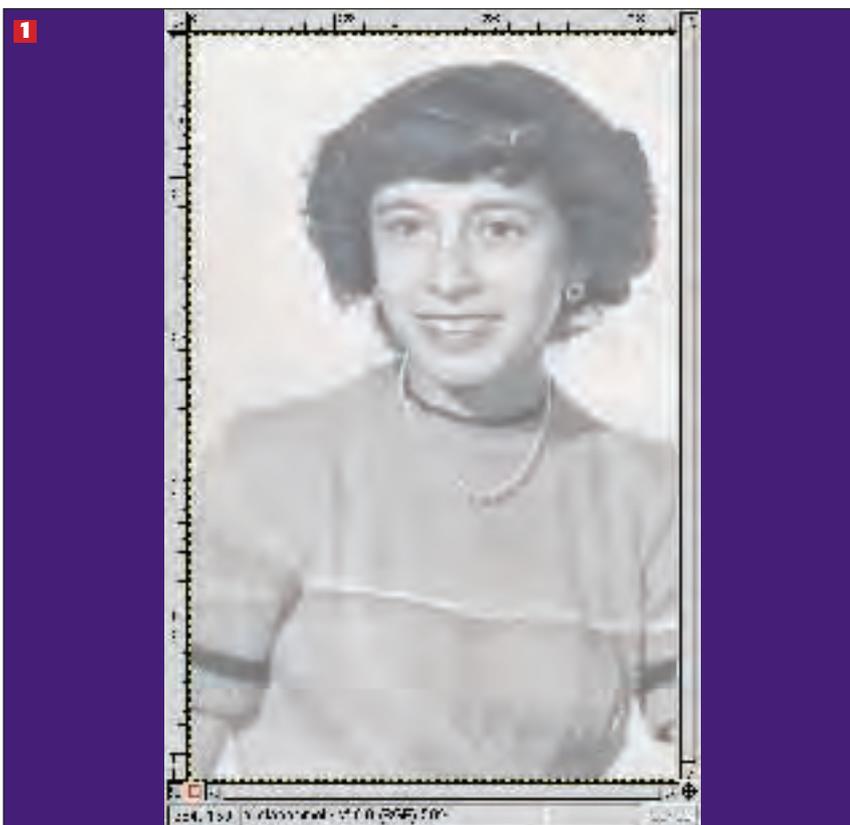
In this tutorial, we'll look in depth at the copy/paste method of fixing heavily damaged images. We'll include discussion on when to use large selections to patch large areas and when to break a blemish into pieces to fix independently. If you're feeling the need to practice the photo restoring techniques you learn here, check out the RetouchPro website (www.retouchpro.com/) and try some of their Challenges that are open to all visitors.

Folds, bends and tears

As you can see, this image has numerous large blemishes. Working on too large an area can cause colour and contrast problems. So we'll break some of these into smaller sections and address them a piece at a time. Use obvious delimiters in the image to choose sections to work on. In this image the area between the arms and midsection offer natural boundaries, as do the hair and bridge of the nose.

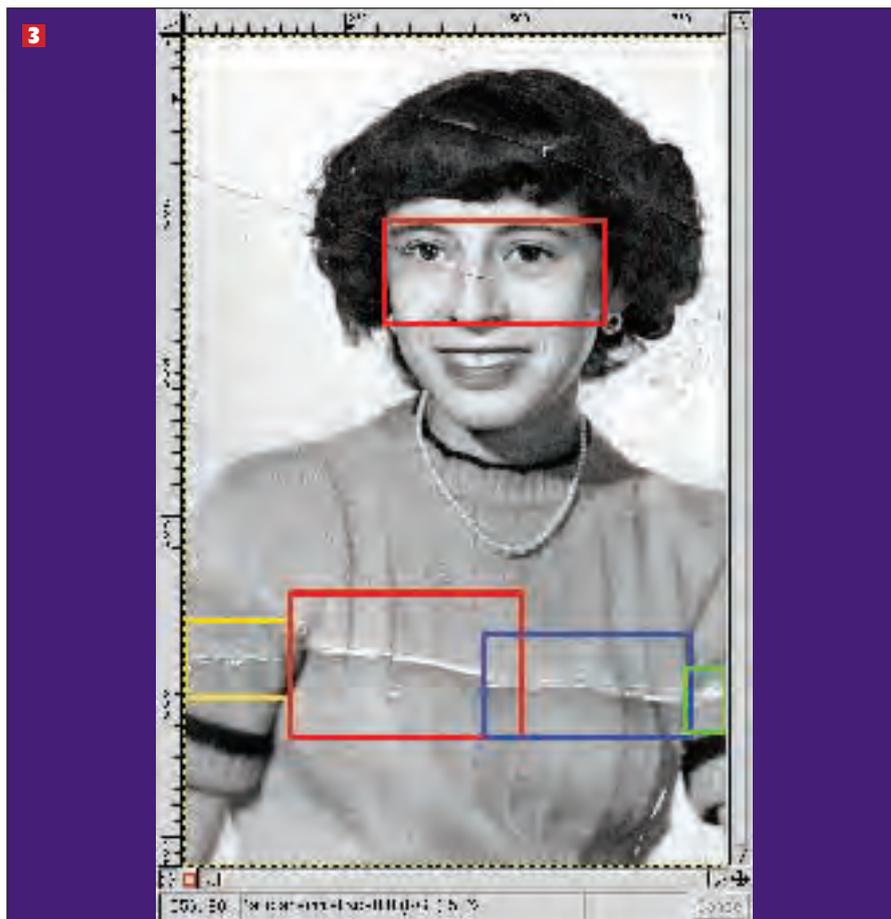
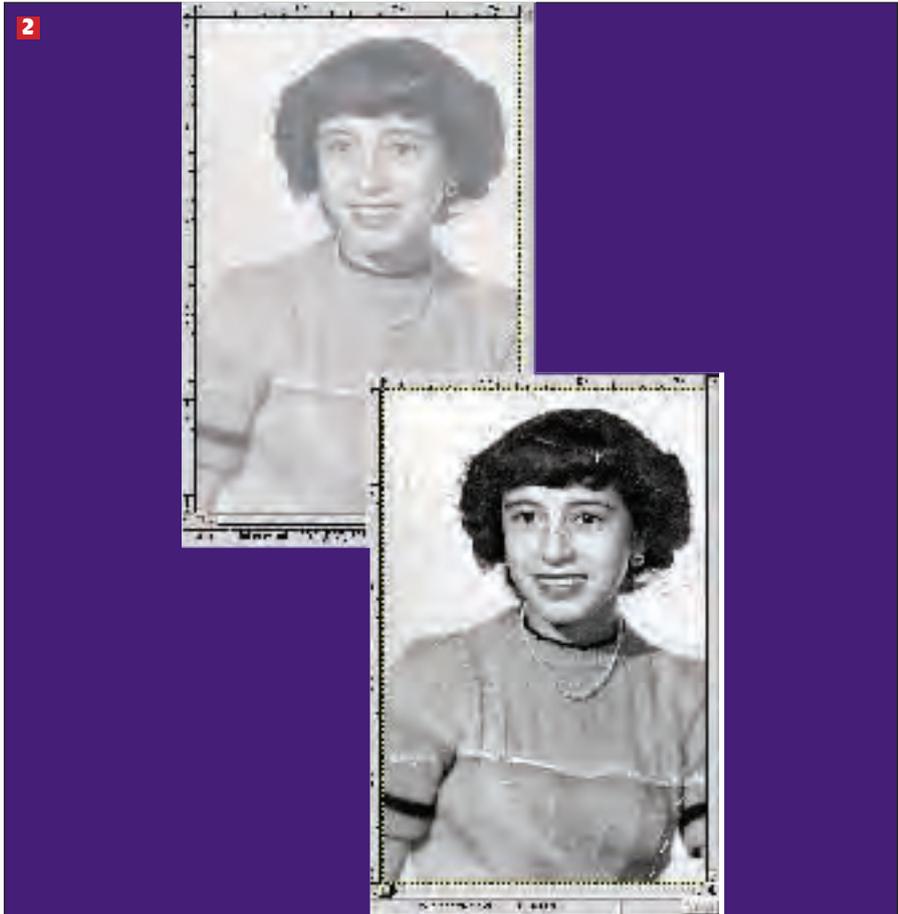
Original scanned image

1 This image was scanned from a 60-year-old photo that had been creased several times, one of which caused the photo to actually tear slightly across the subject's midsection. There are multiple problems to fix here: the faded image needs to have its white and black points corrected, the multiple creases have to be removed where they cross the woman, and the background needs to be cleaned up. There are also barely visible vertical scratches through the midsection, but these can be handled much the same way as the much more visible horizontal scratches and won't be addressed here.



Original/updated comparison

2 We start with some basic image enhancement. We need to fix the white and black points for the scan. We use the Auto-Levels option in the Levels dialog to produce a higher quality image – the contrast is better and the image is clearer. Next we note that the image has taken on some brown tint as it has aged. We clear this by desaturating the image. Finally, as with most scanned images, a little image sharpening is required. We use the Filters>Enhance>Sharpen filter for this process. The UnSharp Mask could also be used although for this image the result would not be significantly better.



Scanner support

Where to get your drivers

SANE
www.mostang.com/sane/
 XSANE
www.xsane.org/
 Vuescan
www.hamrick.com
 Image Scan!
www.epkowa.co.jp/english/index_e.html
 QuiteInsane
<http://quiteinsane.sourceforge.net/index.shtml>
 Epson KOWA
www.epkowa.co.jp/english/linux_e/lsd_e.html
 Linux USB Project
www.linux-usb.org/

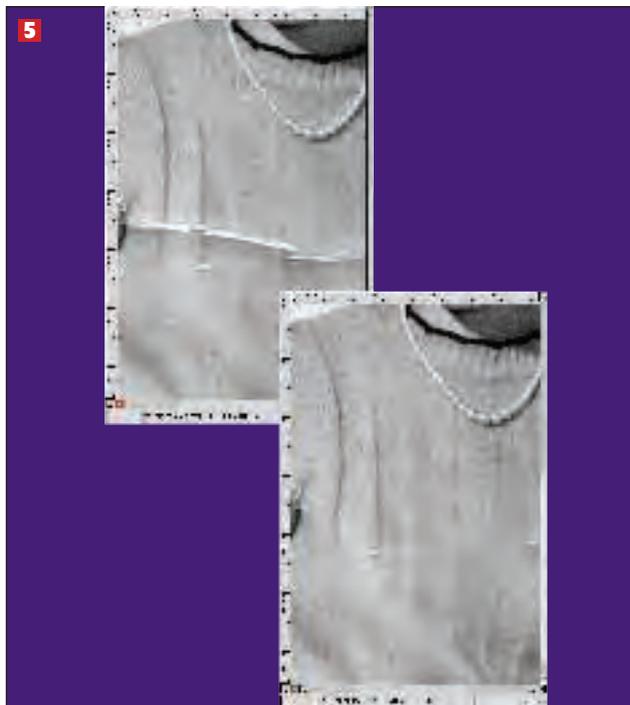
Sections to be worked on

3 There are two main blemishes that fall across the subject: the diagonal fold across face and the horizontal scratch across the midsection. Both will have to be handled by breaking them into multiple and distinct sections and using the same patch layer technique on each. The problem with the midsection scratch is the shadow caused by the scratch itself. If there were no shadow, we could make multiple Fuzzy selects on sections of the scratch, grow them a bit and use this selection to grab a nearby section of the sweater as a patch. But the shadow means we need even a larger selection as a patch so we're better off just using a lasso selection right from the start. >>

TUTORIAL GIMP

◀ Patch positioned over scratch

4 Choose the lasso tool from *The GIMP* Toolbox. Draw a selection around part of the scratch making sure to include the scratch shadow as well. Hold down the **Alt** key (the **Shift** key may also be necessary depending on your Linux desktop configuration) and drag the mouse to move the selection just above the scratch. Use **Ctrl-C** to make a copy of the selection (this will be our patch) and then **Ctrl-V** to paste the selection as a new layer. Position the layer over the scratch. You should be able to line up the pleats in the sweater fairly closely, but perhaps not exactly. We'll fix that in a moment. The original image is tinted red here to show the outline of the patch being applied. Notice that some of the scratch exists in the patch as well. This goes away with the next step – applying a layer mask.



Rip before/after comparison

5 Add a white layer mask to the patch layer. Use the Airbrush to paint black along the edges of the layer inside the mask. This will have the effect of melding the patch into the original image. As long as your selection was large enough you should have a bit of working space in the patch in which to do this melding. If the patch doesn't match the tonal qualities of the scratch area you can adjust its Curves a bit. Only slight modifications to the Curve should be necessary for this image. To line up the pleats more precisely, make sure the layer content is active (not the layer mask) and use the *IWarp* plug-in to move the pleats just a bit. Turn the layer visibility for the patch on and off quickly to see if the patch lines up well with the original image. As you do this it should appear like the pleats are in place and only the scratch is removed.



Sleeves before/after comparison

6 The scratch across both sleeves poses a more difficult problem. There is no unblemished area large enough to cover these scratches so they require multiple smaller patches. In this example we've used six smaller patches to fix the scratch on the left-hand side sleeve. Additionally, the dark area between the sleeve and midsection has a very complex pattern. The patch here will more than likely require Curves adjustment.



Face before/after comparison

7 The face is more difficult than the midsection because of the smooth facial features. Differences between the patch and the face are much easier to notice than between the midsection and its patch. Even so, the facial scratches are best handled by making freehand selections just like the midsection patches. The main difference is that you make much smaller selections, bounded by high contrast lines such as between the face and hair or the bridge of the nose and the shadowed sides of the nose. The eye poses a particularly hard problem because of the complex detail of the image there. In this case, make a selection of the eye, copy and paste it as a patch and then use the Clone tool from the Toolbox to paint out the scratch manually.

Background selection painted with Quick Mask

8 Cleaning up the background is a no-brainer. Make a selection using the Quick Mask and an appropriately sized brush. Note that painting with white will remove the reddish tint of the Quick Mask – that's what you want to happen. You can paint inside curls of hair where the background shows through as well. The painted areas (which are no longer tinted red) will become the selected area when you click on the Quick Mask to Selection button on the bottom left of the Canvas window. Now use the Color Picker tool (the eyedropper icon in the Toolbox) to select a colour from the existing background. Use the Bucket Tool to fill the selection with this colour. Voila! Instant updated background.



Final before/after comparison

9 The results for this image are dramatic. The only visible remnant is a horizontal scratch on the lower right. A closer inspection (not visible in a screen shot) shows a thinner and much longer scratch starting at the left shoulder and running to the bottom of the image. Both of these can be handled with techniques already described.

NEXT MONTH

To create a bit of suspense, we're not telling you what we're covering next month! Get your suggestions for future *GIMP* tutorials in to us at linuxformat@futurenet.co.uk

OUTPUT BUFFERING

Practical PHP Programming

One feature we've consistently been asked to cover in depth is output buffering, so this issue **Paul Hudson** is looking at it in some detail.

Output buffering was introduced in the original release PHP 4, yet for one reason or another it has still yet to become commonly used in websites. This is a great shame, because without output buffering, PHP sends the output of your scripts to your web server as soon as it's ready – this might be line by line or code block by code block.

As you can imagine, the need to send lots of little bits of data is incredibly slow, however much more annoying is the fact that you're restricted in the order you can send data. Output buffering solves this problem by enabling you to buffer up your output and send it to output when you're ready to do so, or even to not send it at all, if you so decide.

We had lots of requests for PHP tutorials on particular topics, and this was the most common. Many people didn't seem to understand the basic OB concepts, so I've broken this tutorial down into bite-size chunks to make each point clear by itself before moving on.

OB Advantages

As most people who work with cookies and other HTTP headers will know, it's often quite a pain to order your output properly. In HTTP, you always need to send header data before content data, which means that if you want to set a cookie half-way through a script, you're in trouble. Luckily, output buffering comes to the rescue by letting you 'send' cookies at any point your script – it stores these cookies separately to the HTML data then sends them together at the end, in the correct order. The bulking together of data also provides quite a performance improvement because there's no longer any need to send it a few kilobytes at a time – PHP stores up all the output of your script until you instruct it to send, at which point all data is sent in one chunk.

The most popular advantage to output buffering is that you can compress content before you send it. Due to the fact that HTML is lots of very simple, repeating text elements, and that normal written text on a web site is also very easy to compress,

compressing your pages can make a big dent in the amount of bandwidth your site (and your visitor!) uses. Because compression requires full knowledge of what it is compressing, you need output buffering.

One of the least-used but most powerful advantages is that output buffers are stackable, meaning that you can have several buffers working on top of each other, allowing you to build your output up over multiple buffers.

Performance considerations

If you're not using content compression, output buffering is very unlikely to affect the speed of your web server by any great amount – if anything, it should help it serve pages *faster* because of the optimised data sending. Content compression does take up a little CPU time on both the server and on clients visiting your site, but it's pretty small.

On the up-side, content compression should decrease the amount of bandwidth you use by 40-60%, which means your server will spend less time sending data across the network. The compression level you achieve depends entirely on the kind of content you serve up – if you have lots of pictures, which content compression won't affect, your compression level will be lower; if you're sending lots of XML, which is a naturally repeating format that is very easy to compress, your compression level will be much higher. It's important to remember that only the output of your PHP script will be compressed – images, CSS files *etc* are all served as normal.

Getting started

You can enable output buffering in one of two ways, one of which seems easier at first but is likely to cause hassle in the long term. The 'easy' option is to edit your php.ini file to enable output buffering for all scripts – this might sound great, but it will mean your scripts will break on other PHP installations, and also means you have no way to not have output buffering for a script. The second option, which is much smarter, is to use the set of output buffering function calls on a script-by-script basis.

The `ob_start()` function is used to create a new output buffer, and you can immediately start writing to it by printing out content as normal. Once you have a buffer open, there are two ways to close it: `ob_end_flush()` and `ob_end_clean()`, both of which end the buffer, but do so in slightly different ways. The former ends the buffer and sends all data to output, and the latter ends the buffer without sending it to output, effectively wiping out any information you saved in there. Every piece of text outputted while an output buffer is open is placed into that buffer as opposed to being sent to output. Consider the following script:

```
<?php
ob_start();
print "In first buffer!\n";
ob_end_flush();
ob_start();
print "In second buffer!\n";
ob_end_clean();
ob_start();
print "In third buffer!\n";
?>
```

That script will output **"In first buffer"** because the first text is placed into a buffer then flushed with `ob_end_flush()`. The **"In second buffer"** won't be printed out, though, because it's placed into a buffer which is cleaned using `ob_end_clean()` and

Make your mark

Brainstorms 'R' Us

Would you like to get your name in the mag and learn about stuff you're most interested in?

We're always looking out for ideas for new *Linux Format* PHP tutorials, and where better to look than to you, our readers? If, while reading past issues of *LXF's* PHP tutorials, you've thought "I wish they'd covered XYZ in more depth...", or "I really want to know how to use...", then now's the time to get your voice heard!

Send an email to paul.hudson@futurenet.co.uk with your ideas – all the good suggestions that you send in will be covered in future issues. So far, the topics we have covered in some depth include MySQL, XML, CLI, GUIs, media generation, templates, and more.

If you're short of ideas, you're certainly welcome to write in or post on the forums at www.linuxformat.co.uk with comments – we're passionate about improving the overall quality of our tutorials!

not sent to output. Finally, the script will print out **"In third buffer"** because PHP automatically flushes open output buffers when it reaches the end of a script.

Stacking buffers

The functions `ob_end_flush()` and `ob_end_clean()` are complemented by `ob_flush()` and `ob_clean()` – these do the same jobs as their longer cousins, with the difference that they don't end the output buffer. Instead, these functions send the content to output or clean the buffer (respectively), leaving it open for more text. We'll be looking at how you can use these functions to re-use your buffers later on, but for now it's important to understand that you can stack buffers up upon each other to make them even more useful.

Consider the following script:

```
<?php
ob_start();
print "In first buffer!\n";
ob_start();
print "In second buffer!\n";
ob_clean();
?>
```

In that script, we call `ob_start()` twice without closing either of the buffers, and so the end result is that **"In first buffer"** is printed out by itself. If you thought that **"In second buffer"** would be printed out too or that neither lines of text would appear, you haven't grasped quite how buffer stacking works!

The first buffer is started and filled with **"In first buffer"**, then a second buffer is started on top of the first buffer, leaving the first buffer still intact and containing **"In first buffer"**. At this point, we can no longer write to the first buffer, because the second buffer is top of the stack. The new buffer is filled with **"Hello second"**, and finally `ob_clean()` is called, wiping the second buffer, *but leaving the first one intact*.

Flushing stacked buffers

When you stack your output buffers up, data you flush is moved up one level in the stack as opposed to being sent directly to output. This makes more sense with some code, so here you go:

```
<?php
ob_start();
print "In first buffer!\n";
ob_start();
print "In second buffer!\n";
```



```

<< ob_end_flush();
    print "In first buffer\n";
    ob_end_flush();
?>

```

That script will output the following:

```

In first buffer
In second buffer
In first buffer

```

What happens there is that the second buffer gets flushed into the first buffer where it left off, as opposed to directly to output – it literally gets copied into the parent buffer. The first buffer then gets **“In first buffer”** added to it, then flushed to output. Take a look at this following script:

```

<?php
ob_start();
print "In first buffer\n";
ob_start();
print "In second buffer\n";
ob_end_flush();
print "In first buffer\n";
ob_end_clean();
?>

```

It appears to be the same as the previous script, with the only difference being the last line – **ob_end_clean()** is used rather than **ob_end_flush()**. This time the output is nothing at all, because the second buffer gets flushed into the first buffer, then the first buffer gets cleaned, which means the clients receives none of the text.

As long as you keep in mind that output buffers are stacked up like blocks, which means you can't write to any one of the stack of buffers, this functionality will work in your favour. Using this method it's very easy to progressively build up your content by opening up new buffers as needed, flushing in content to a parent buffer as you go.

Reusing buffers

Given that **ob_flush()** and **ob_clean()** leave the current output buffer open for further writing, there's a potentially big performance boost just waiting to be taken advantage of – by not closing and re-opening buffers all the time, this next script could be rewritten...

```

<?php
ob_start();
print "In first buffer!\n";
ob_end_flush();
ob_start();
print "In second buffer!\n";
ob_end_clean();
ob_start();
print "In third buffer!\n";
?>

```

...like this...

```

<?php
ob_start();
print "In first buffer!\n";
ob_flush();
print "In second buffer!\n";
ob_clean();
print "In third buffer!\n";
?>

```

In the new script, the buffer is first flushed and left open, then cleaned and still left open, until finally being automatically closed and flushed by PHP as the script terminates. By not needing to create and end output buffers as the script executes, thereby reusing the same buffer each time, that script executes about 60% faster – this is a substantial difference, as I'm sure you'll agree.

Reading buffers

While writing to and flushing buffers is a boon by itself, you can also *read back* the contents of output buffers, effectively receiving a copy of all the output it holds. This clever functionality is contained in one simple call to **ob_get_contents()**, which takes no parameters and returns a string of all the content it contains. Reading your output back from a buffer is more useful than you might at first think, but it does take a little experimenting to get quite right.

Last issue, for example, we used output buffering and **ob_get_contents()** to write a static page cache – the modified date of the PHP script was compared against the modified date of the cached page, and if the script was newer, it would execute and output its content into an output buffer, which was then retrieved and written to a file.

The key advantage to retrieving output buffering is that you can make one script do many things with almost no change. For example, if you have a script that tracks the location of a package while it's being shipped around the world, the default configuration might have it send its data directly to output for web browsers. However, by using output buffering it literally is a tiny change to make that same script send its output to email, or to an SMS number – the possibilities are endless.

Combining reading output buffering with flushing means that you can save your output to a buffer, read it back in, pass it through various functions to alter the data, then send it back to output – you really have much, much more flexibility, and there are many clever ways you can take advantage of this.

Other OB functions

There are two utility functions that give you information on your current output buffering situation, and these are **ob_get_level()** and **ob_get_length()**. The **ob_get_level()** function is particularly useful as it tells you the buffer stack level you're at – literally how many buffers you have open currently. By default this will return 0 because you have no buffers open, but this number increases as you add more buffers. **ob_get_level()** function is particularly helpful if you want to recursively work with and close open buffers, because you can loop down from **ob_get_level()** to 0.

On the other hand, there's **ob_get_length()**, which returns the size in bytes of the current output buffer. Note that this is not the total length of all buffers, but only the length of the current buffer – you need to use this in combination with **ob_get_level()** to get the total buffer lengths while flushing.

Compressing output

When visitors come to your site, the HTTP request they send for a page also includes a lot of other information about that visitor. For example, it sends the name of the web browser they are using, the last page they were at, and what kind of content encoding they can accept. The content encoding is what we're interested in, because browsers that support compressed (*gzipped*) HTML say so in the HTTP request, which means that the web server can *gzip* the content before it sends it. A key feature of this system is that if a client *doesn't* say that it supports

gzip encoding, the web server sends back plain text – this process is all entirely transparent to the user.

Output compression, being just *gzip* wrapped up nicely, requires that you have all the content ready before you compress it. Naturally this is very close to output buffering, which collects all its content up in buffers before sending it out. As such, it's an easy jump from output buffering to content compression – you store your data up in a buffer, compress it at the last moment then send it out.

Passing the parameter `ob_gzhandler` to `ob_start()` enables output buffering compression – PHP automatically takes care of checking whether the client supports it or not, and only sends compressed output if it is supported. By calling `ob_start()` with `ob_gzhandler`, you're effectively saying "if compression is supported, send content compressed; otherwise, send plain text".

At the client end, the compressed content is automatically uncompressed and displayed normally, because the process is transparent. If you want to check whether your output compression is worked, you need to telnet into your server. First, create a very basic script like the one below:

```
<?php
ob_start('ob_gzhandler');
print "Goodbye, Perl!\n";
?>
```

Remember that PHP automatically closes and flushes open output buffers, so the above script simply outputs the compressed text, "**Goodbye, Perl!**". Save the script into your public HTML directory as `gztest.php`. To check that the `ob_gzhandler` parameter is working, you need to telnet into your server on port 80, so, from a console window, enter this command:

```
telnet <your server> 80
```

This will connect to your web server on the HTTP port – you now get to pretend to be a web browser. Once you're connected, enter `GET /gztest.php HTTP/1.0` and press enter twice. This forms your complete HTTP request, and you should get your response quite quickly – this should be several lines of HTTP headers, followed by the content of the page, "**Goodbye, Perl!**".

Here's what I got below:

```
paul@hud-lxf:~$ telnet localhost 80
Trying 127.0.0.1...
Connected to localhost.
Escape character is '^]'.
GET /gztest.php HTTP/1.0

HTTP/1.1 200 OK
Date: Mon, 04 Aug 2003 10:43:31 GMT
Server: Apache/1.3.28 (Unix) PHP/4.3.3
X-Powered-By: PHP/4.3.3
Connection: close
Content-Type: text/html

Goodbye, Perl!
Connection closed by foreign host.
```

After the usual collection of HTTP headers coming back in the response, "**Goodbye, Perl!**" is there – it's not compressed though. The reason for this is because, as mentioned, web servers will only send compressed output if browsers say that they support compression. So, to get compressed content back, we need to mimic a browser that supports compression. Open up the telnet connection again with the same command as last time, but this time enter `GET /ob.php HTTP/1.0`, press **Enter** *once*, then type

```
ACCEPT-ENCODING: gzip
```

and press **Enter** twice. Here's what I got this time around:

```
paul@hud-lxf:~$ telnet localhost 80
Trying 127.0.0.1...
Connected to localhost.
Escape character is '^]'.
GET /gztest.php HTTP/1.0
ACCEPT-ENCODING: gzip

HTTP/1.1 200 OK
Date: Mon, 04 Aug 2003 10:43:31 GMT
Server: Apache/1.3.28 (Unix) PHP/4.3.3
X-Powered-By: PHP/4.3.3
Content-Encoding: gzip
Vary: Accept-Encoding
Connection: close
Content-Type: text/html

œœOI(tm)L'H- Q%œ,I.
Connection closed by foreign host.
```

Note that our content has now been compressed – it's almost certainly not going to come out properly printed in the magazine, but if you try it yourself you'll see that it outputs various obscure ASCII characters instead of "**Goodbye, Perl!**". You'll also see that the compressed version is actually *longer* than the uncompressed version – this is because compressing very small amounts of text rarely has any gain, and usually actually works out worse. However, it's rare that your site would contain such a small amount of text – once you reach about 100 bytes, the compression will start to work in your favour.

As you've seen, the same PHP code can send two different pages depending on whether the client signals that it supports content compression. As such, there really is little reason why you shouldn't use compression unless you're particularly short of CPU time and have bandwidth to spare.

Conclusion

Hopefully this excursion into the world of output buffering has cleared up any questions you had, and also hopefully expanded your horizons a little. Output buffering still hasn't got a big following in the developer community, which is a shame – it's powerful, fast, and opens up so many new possibilities for programming that would simply not be possible otherwise.

For many, output buffering is a 'neat feature' that's rarely used, simply because they often don't understand quite how big an impact it can have. So, I encourage you to throw your inhibitions to the wind and give it a try – I'm almost certain you won't regret it! .

PHP 4.3.3 released!

PHP have released new versions recently. If you are running PHP on a production server, or anywhere accessible from the Internet, the developers of PHP strongly recommend that you upgrade to 4.3.3. It contains a large number of bug fixes, some of which are security related. If you prefer to play with bleeding-edge code and are not using PHP in a mission-critical system, you may want to try the first beta release of PHP 5.0. This won't be suitable for a production system – beta releases *are supposed* to have bugs – but if you feel confident enough to try the latest version, and help development by reporting any bugs you discover, please give it a go! [LXF](#)

NEXT MONTH

We're going to be looking at how PHP's socket functions work, and how they allow you to make a web server entirely in PHP.

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



Snails' pace box

Q I recently installed Red Hat 8.0 which I got from a friend as I had heard that Linux was good on older smaller PCs (mine is a custom-built AMD Athlon 500 with 64MB RAM) and I so want to ditch Windows, or at least not be as dependant on it.

The question is, why does Linux take literally two or three minutes to open any application. Indeed it takes five minutes to boot and a good few minutes to start the desktop (either GNOME or Blackbox which I downloaded).

Windows flies on my machine but this really creeps along at an infuriating pace. As I had heard so many good things about Linux I am wondering if it was me that has done something wrong?

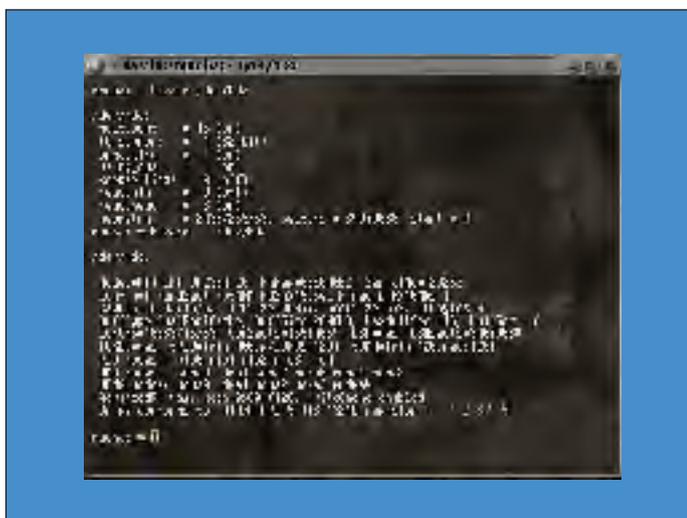
I have only installed the standard installation package but really, 2 or 3 minutes just to open a console or note pad and 15-20 seconds for the start menu to appear not to mention the thrashing my hard drive is subjected to. This also happens when closing programs. The program stays there for a good minute or so before disappearing?!

Does anyone know how to speed this thing up?

Andrew, from LXF forums

A You didn't say what you were running, so we can only speculate as to why your system is running slowly. There are three main reasons why Linux can run slowly; insufficient CPU time, lack of DMA support for storage devices, or lack of RAM causing swapping.

If you are running something like KDE or GNOME under 64MB RAM, you'll probably find that you start to swap pretty quickly. Of course, if it takes five minutes to boot up, then there is probably something else wrong with it.



Using **hdparm** to tune the IDE drives on a system can boost throughput, but ensure that there is support for the IDE chipset in the kernel.

Checking DMA on your devices is a good first step, done with **hdparm**:

```
# hdparm /dev/hda
/dev/hda:
multcount   = 16 (on)
IO_support  = 1 (32-bit)
unmaskirq   = 1 (on)
using_dma   = 1 (on)
keepsettings = 0 (off)
readonly    = 0 (off)
readahead   = 8 (on)
geometry    = 2432/255/63,
sectors     = 39070080, start = 0
# hdparm -i /dev/hda
/dev/hda:
Model=HITACHI_DK23BA-20,
FwRev=00E0A0E2,
SerialNo=12J26Z
Config={ HardSect NotMFM
HdSw>15uSec Fixed DTR>10Mbs }
RawCHS=16383/16/63,
TrkSize=36477, SectSize=579,
ECCbytes=4
BuffType=DualPortCache,
BuffSize=2048KB, MaxMultSect=16,
MultSect=16
```

```
CurCHS=16383/16/63,
CurSects=16514064, LBA=yes,
LBAsects=39070080
```

```
IORDY=yes,
tPIO={min:400,w/IORDY:120},
tDMA={min:120,rec:120}
PIO modes: pio0 pio1 pio2 pio3
pio4
```

```
DMA modes: sdma0 sdma1 sdma2
mdma0 mdma1 mdma2
UDMA modes: udma0 udma1
udma2 udma3 *udma4
AdvancedPM=yes: mode=0x80
(128) WriteCache=enabled
```

```
Drive conforms to: ATA/ATAPI-5 T13
1321D revision 1: 1 2 3 4 5
```

Assuming that the IDE chipset is supported by the kernel, one can enable DMA and set the appropriate mode using the following:

```
# hdparm -d1 -X udma4
```

where **udma4** is one of the supported UDMA modes from the **hdparm** output. You may also want to run a **ps ax** to find out exactly what is running, so that you can disable anything you don't need.

A QUICK REFERENCE TO: Installing Software

Some Linux distributions use packages, some using source code and others using a combination of the two, maintaining a system can sometimes become quite a chore. Not to mention the fact that there are half-a-dozen different distros out there using RPMs, all of which require a RPM specific to the distro. We're going to look at installing software using every method known to mankind, hopefully covering all package systems used on common Linux distros.

Install from sources

While all packages don't work correctly on all distributions, sources will work on everything from Red Hat to Debian, even on non-Linux operating systems. Not all sources compile in the same way, although there is a general rule of thumb to follow. If all else fails, there is a usually a README or INSTALL file to read describing the process to compile, install and configure the software. Quite how well these files are written is anyone's guess, since more than often developers spend more time writing code than describing how to use it.

For the vast majority of software, the first step is the **configure** script. This will perform numerous tests on the system, looking for libraries, the C compiler, required headers and generally validating specific information required by the software to compile, much less actually run. **configure** is run by using **cd** to enter the source directory then doing:

```
$ ./configure
```

The **configure** script will generally have a wide range of options to influence the generated binaries, so this is where we tell it what libraries we want it to link against, what capabilities it should be compiled with, plus pointing it to any libraries or headers located in a non-standard location.

If **configure** decides that everything is good, then we can proceed to build and install the software. Everything from this

point on is done with **make**:

```
$ make
$ make test
```

The actual installation must be performed as root, if the software is not to be installed in a user's home directory. Using sources allows a user without root privileges to compile and install software in their own home directory. Packaging systems require root privileges to read and modify the package information located in **/var**, so often source is the only option to get something installed without obtaining root.

Following a **make install** as root, the software is installed and available. As default, it will be installed in **/usr/local/**, so any binaries will live in **/usr/local/bin** and shared libraries in **/usr/local/lib**. On systems where source code is being used frequently, it's useful to add **/usr/local/bin** to any **\$PATH** variables in **/etc/profile** or **/etc/bash_profile**, plus **/usr/local/lib** should be added to **/etc/ld.so.conf** as it is not a standard search path.

It's worth remembering that source installations won't do a whole lot beyond actually plonking default configuration files in the right place, so manual editing is required plus one will likely have to write a custom **init.d** script or locate one on the Internet, as **init.d** scripts are specific to each distribution.

RPMs

Most distributions use the Red Hat Package Manager, or RPM. Red Hat, Mandrake, SuSE and plenty of others allow the administrator to install software using RPMs, although it's important to install the RPM specific for the distribution, as configuration information and scripts are generally quite different between them all.

Installing an RPM is fairly basic, as we just need to do:

```
# rpm -Uvh traceroute-1.4a5-24.6x.i386.rpm
```

After a short while, **rpm** will validate the RPM, install it and



alien allows us to convert Red Hat packages to Debian, Slackware and Stampede formats. Get it from <http://kitenet.net/programs/alien/>

proceed to do any configuration it needs to actually make the installed software do something useful. Of course, as anyone who has used RPMs on systems will know, nine times out of ten, a few dependencies work their way out of the background to bite us.

Usually it's now time to head over to <http://rpmfind.net/>, figure out what it's missing and try to install what it is asking for. Should the dependency require yet another package, back we go to <http://rpmfind.net/> and continue the process.

Once the RPM is installed, it will be included in the package list on the system and everything that is needed, including basic configuration, **init.d** scripts and any modifications required to the system, such as adding new users or changing permissions. See Hoyt Duff's two-part tutorial in *LXF44* and *45* for info about alternatives to the Red Hat Package manager.

Debs

Debian users have the **deb** package format, which are usually installed using the rather slick **apt-get** tool. Starting with an **apt-get update** to grab the latest package listings from the servers, we can install any package we want which is included in the list.

Searching for a package is done using **apt-cache** and we just do:

```
# apt-cache search php4
```

Every package with something referring to **php4** in its name or description will be spat out, so pushing it all through **grep** to limit the output is usually not a bad idea. The desired package can then be installed using **apt-get**:

```
# apt-get install php4-cgi
```

apt will download the desired package, including dependencies, then proceed to install and configure each one in turn. Most Debian packages install without any interaction, but some ask for run time options, to enable specific features. One can reconfigure the package later by doing:

```
# dpkg-reconfigure php4-cgi
```

Convert packages

More often than not, software is distributed in source form, along with the preferred package format of the developer. Unfortunately, this means that alternative packages must be created by a third-party, so the lead-time on some things is a little excessive. However, the **alien** tool allows us to change a package from one format to another. If we have a RPM, we can create a **.deb** from it so it can be installed on Debian systems. The conversion process is not as smooth as if the package was designed for the distribution, so scripts may have to be tweaked to ensure that everything works as it should.



« Vintage Wine

Q I've been reading *LXF* for about a year now and find it gives a comprehensive coverage of things Linux. Do you know where I can obtain a version of Winelib for Solaris for Intel? There seems to be a version of Wine that is available at www.ecn.purdue.edu/~laird/WINE/, but I can't seem to find a version of Winelib.

Michael Byrne, via email

A The version of Wine you refer to is obscenely old, predating the creation of *libwine*. You should be able to compile the current Wine source tree under Solaris, and refer to the Wine mailing lists for any compilation failures. We couldn't find any pre-built Wine binaries for Solaris x86, although the Wine lists may yield a contact who either has pre-built Wine binaries or distributes them.

Modem madness

Q As a complete newcomer to Linux (and new subscriber to *LXF*), and having a certain amount of difficulty with my recent installation of SUSE 8.2, I am very pleased to have articles like the *Beginners' Guide to Linux* in *LXF43*. However, at the moment, I have managed to load SUSE happily on to the same hard drive as my WinXP and can dual boot without any problems. I am slowly feeling my way round the new system and have configured my Epson scanner/printer and separate printer.

BUT – the Internet is altogether another matter! I am aware that Winmodems are 'difficult', not to say almost impossible to configure, but I had an external USB/serial port Diva 852 ISDN T/A working perfectly under XP. I also had a Xircom PC Card 56K modem working fine in a PCI PCMCIA bus card under Windows XP. Under Linux I can get absolutely nothing to even 'blink' with either! I started with the Diva connected on the serial port as I believed Linux would recognise any external serial modem – it doesn't, or certainly not automatically. It does report that I have two AT modems installed, but that's all. My efforts at configuring them do not seem to work. I tried the USB connection – hot-plugging – and Linux



Need an external modem that works with Linux? Turn to page 15's *Reader Tips* where a couple are suggested.

immediately reported it and I followed the configuration set-up and it appeared to install as an external USB Diva 852 ISDN T/A. But I cannot get it to dial out. I have asked SUSE Helpdesk and they say the T/A is not a 'real' serial modem/T/A, but in fact another Winmodem!! Hey- Ho!

So – and sorry to be a little long-winded – 2 quick questions:

1 Can/should I be able to get either of my existing 'modems' to work with Linux?

2 The article in *LXF43* mentioned that a suitable PCI modem for Linux should be available for about £15 – and 'more on this in the future'! Can I please ask to 'jump the gun' here, and ask for a model/version number that I can easily get hold of to work with my ISDN line and Linux? This, I suspect, will be the best way to go anyway, but I really would like to persevere with Linux and without an Internet connection it will be hopeless to migrate across from XP in any sensible form.

Ian Leckenby, via email

A Many USB 'modems' are indeed Winmodems, although the manufacturers rarely admit to it. Unless you have found the USB modem or terminal adaptor on linux-usb.org, it's generally best to just avoid them. However, if you can connect it via a serial port, then it should do something. At the very least you should be able to fire AT commands back and forth using something like *minicom*. There are also many PCMCIA Winmodems out there, so if it's not supported by *pcmcia-cs*, it's probably a Winmodem. There is a fairly complete list of what does, and does not, function with the Linux PCMCIA drivers at <http://pcmcia-cs.sf.net>. For both devices, the chances of getting them to work are somewhere between slim and non-existent, although

linmodems.org may provide you with some useful pointers.

Eicon has a number of Diva PCI terminal adaptors which work well with the Linux ISDN kernel drivers. A complete list can be found at www.isdn4linux.de/faq/i4faq-5.html, which are supported using the ISDN4Linux kernel modules.

Waiting for a bus

Q I'm pretty much a newbie, having only been running Red Hat 8 at home since January this year, but I've found it pretty easy to pick up. I seem to find enough information within your pages each month to keep me reading (and tweaking) until the next issue comes out.

I'm thinking of adding a second hard disk to my desktop system and making it double as a file/print server for my home network (well my laptop). I have a few questions about "best practice" which I hope you can answer.

My current setup has the hard disk installed as hda. The second IDE controller has a DVD and an IDE/ATAPI CDRW. I want to install the new hard drive as the secondary IDE master, reconnecting the DVD as hdb, but I have heard that this will slow both IDE buses to ATA33. Is this true, or will I get better throughput with the hard-drives on separate controllers?

Secondly, I already have a 12GB partition mounted as /data. Is it OK to create my new mount points within this partition, or would it be better to unmount the current partition, create my new mount points there, and then remount the old data partition a level further down? Is there a recommended max partition size for efficiency?

Finally, as my laptop will dual boot between Windows and Linux, I

know I will need to run *Samba*. Would I gain any benefits from simultaneously sharing the filesystem with *NFS*, and the printer with *CUPS* for the times when the laptop is running Linux?

David W Haggett, via email

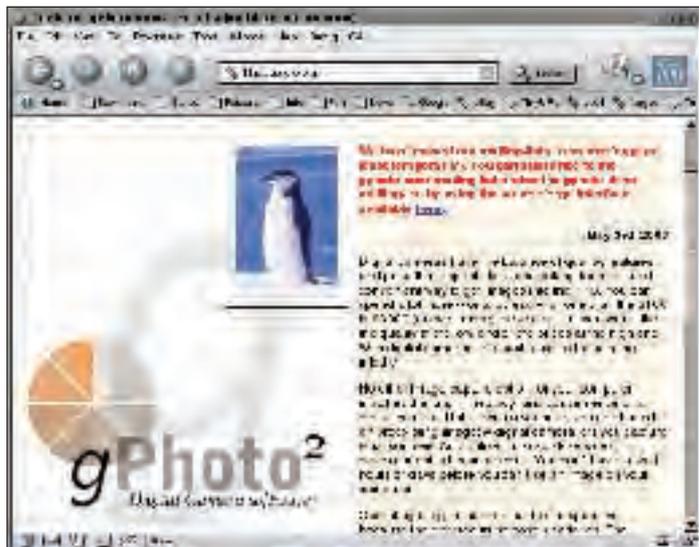
A When you connect two devices to an IDE bus, you are limited to the throughput of the slowest device. It makes far more sense to keep the primary IDE controller for hard drives and use the secondary for the DVD and CD-RW. Having the two hard drives on the same controller is not usually a big problem, unless you're running some sort of RAID where two writes may be occurring to the two drives at the same time.

You can mount the partitions from the new drive anywhere you want, although it's a good idea to go over your existing directory structure using `du -sh <dir>` to figure out how much disk space is being used. There is really no efficiency advantage by arranging your partitions a particular way, and the exact usage of your system will depend where you actually mount the new filesystems.

Sharing filesystems with *NFS* when you're using Linux has a number of advantages, including the fact that UID/GID information is shared correctly. It's usually much easier to share printers with *CUPS* or *lpd* between Linux boxes than to mess with *smbclient* and figure out how to conjour up a configuration to make it work well. It does mean that you'll have to maintain *NFS*, *CUPS* and *Samba* all at the same time, but it may make life a little easier initially.

USB controller

Q I have been desperately trying get one of two Linux distros to work, these being SuSE 8.0 pro and Mandrake 9.1.



gphoto2 supports a wide range of digital cameras, but does not require any kernel support beyond basic USB functionality.

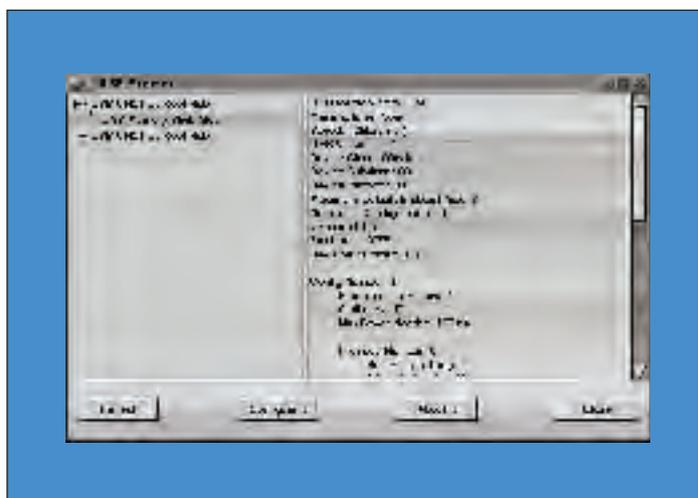
The former I purchased two years ago and I'm itching to try it out but have been put off by the difficulties in customising my installation to work with my existing USB peripherals namely the following:

- Scanner – Canon N650U
- Camera – FinePix 6800Zoom
- Modem – Speedtouch 330 ADSL

I've visited every forum to ask every conceivable question but have been frustrated at every turn.

This is what I see when looking at the hardware listing in the Control Center/USB view:

USB Controllers:
7001
Vendor Silicon Integrated Systems
Bus PCI
Bus ID 1039:7001
Location on Bus 0:2:2
Module USB_OHCI



Inspecting USB devices using usbview shows what is connected to a system and shows obvious problems with devices.

Media Class SERIAL_USB
USB OHCI Root Hub:
Bus USB
Module unknown (This only in Mandrake 9.1)

From what I understand, I need to load a particular module but I've tried this both with *xconfig* and even more drastically recompiling the kernel (used a reference book step-by-step, and though not entirely happy I got it right).

Would the above indicate that there is a compatibility issue with my PC's USB root hub? I'm determined not to be one of those Linux newbies who, after an initial wave of enthusiasm, is forced back to Windows because it is so darn difficult to configure.

Mark Olivier, via email

A It looks like the kernel found your USB controller, so it's just a case of going through each of the devices one by one to provide kernel support for it. Some USBs are accessed via kernel modules, others are used directly via user-space processes and require no specific kernel modules.

The digital camera is the easiest, as it is supported by *gphoto2* (<http://gphoto2.org>) without any additional kernel modules. You can simply run

```
gphoto2 -P -R
```

to have it autodetect your camera and download the photos, although there are a number of different GUI front-ends to *gphoto2*.

The Speedtouch ADSL modem is also supported by Linux. Some documentation on getting it to work with Linux can be found at www.linux-usb.org/SpeedTouch/.

Your scanner is also well supported by Sane (<http://panda.mostang.com/sane/>). You should be able to find a package of *Sane* for Mandrake, otherwise you will have to manually compile it.

Graphics card hard

Q I have a strange problem with Mandrake 9.1 that came with the Linux Format 41, and hope you can help me. It seems my computer doesn't like anything else but Win XP. My Comp. Config. is the following:

OS Name	Microsoft Windows XP Professional
Version	5.1.2600 Build 2600
OS Manufacturer	Microsoft Corporation
System Manufacturer	GBT
System Model	AWRDACPI
System Type	X86-based PC
Processor	x86 Family 15 Model 2 Stepping 7 GenuineIntel ~2642 Mhz
BIOS Version/Date	Award Software International, Inc. F3, 26/12/2002
SMBIOS Version	2.3
Windows Directory	C:\WINDOWS
System Directory	C:\WINDOWS\System32
Boot Device	\Device\HarddiskVolume1
Locale	United States
Time Zone	W. Australia Standard Time
Total Physical Memory	1,024.00 MB
Available Physical Memory	776.76 MB
Total Virtual Memory	2.65 GB

Available Virtual Memory	2.26 GB
Page File Space	1.65 GB
Page File	C:\pagefile.sys

I have 2 hard drives, one of which has XP installed, and the other one was meant for Linux. The video card is a Gigabyte Maya Radeon 9000 Pro, and I also have a Asus Burner, and Toshiba DVD.

And here is my problem. When I boot from the Mandrake disc, the splash screen appears, and after pressing **Enter**, the screen goes black and stays black, and I have no option but to turn off the computer. And here is the funny bit – I had no problems at all installing Mandrake 9.1 on my son's computer, who is running XP as well. I'd like to install Mandrake on the second HDD and would appreciate any suggestions.

Hans, via email

A It's more than likely that Mandrake 9.1 does not like your video card, so can't initialise it properly when it tries to switch to GUI mode during the installation. You can do a text-based install by passing **text** at the install prompt, which will avoid it trying to figure out your video card.

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: lxf.answers@futurenet.co.uk

ANSWERS



Mailserver apps

Q Whilst setting up our home network, I was faced with the dilemma of setting up a local mail server – whilst not totally necessary it would be nice to have. My question is, what would be your recommendation for mailserver software, and would it be possible to set it up in such a way that when someone sends a mail the mail server (or probably to be more correct SMTP server) would check to see if the to address was an internal or external address, and upon this information either route the mail internally or forward it on to the big wide world?

Daryl, via email

A *Sendmail* as a MTA can certainly handle mail routing between local domains and the public Internet. Depending on which version of *sendmail* and which Linux distribution you are using, you'll need to add your local domain or domains to one of the file `/etc/mail/local-host-names` (on some very old versions of *sendmail* this is `/etc/mail/sendmail.cw`). The format of this file is one line per domain eg:

```
localhost
domain.com
otherdomain.com
```

For any changes to take affect here you'll need to restart *sendmail*.

Reverse lookups

Q Under DNS, I am quite happy setting up zones for the resolution of host names into IP addresses, and the delegation of sub-domains, but I'm struggling to get reverse lookups to work, converting an IP address back into a hostname.

I can see how the `x.y.in-addr.arpa` domain is used for reverse lookups, how it can delegate complete class A, B and C address ranges, but my ISP only allocates me a partial class C address, comprising 32 addresses. How do I get this partial subnet delegated to me so that I can

control the lookups of addresses in this range?

The only solutions I can see are for the whole class C range to be delegated to me (but surely that would let me control the lookups of addresses outside my range as well as my own), or to let the ISP control the whole zone, but that doesn't give me the flexibility I want.

Mark, via email

Unfortunately it is not possible to delegate control of parts of class C subnet for reverse DNS lookups. The only recommendation I can make is to contact your ISP and ask them to set the PTR records on their name servers. Most ISPs today will not delegate control over entire class C subnets to their customers. The only way around this is to approach your local IP address authority (RIPE in Europe, ARIN in the US) and ask them to assign you a full class C. This probably won't be very easy and depends on the size of your organisation and your IP address usage. If you're like most of us you'll probably need to trust your ISP with this responsibility.

Migrating to Linux

Q Our company used to be a complete Microsoft web development company.

Recently we've had to start supporting Linux too due to phenomenal demand. I've been tasked with taking care of the migration of certain sites from the Windows 2000 server to the newly installed SUSE Linux Professional server. We had some consultants come in to do the installation.

The problem I have is that whenever we move files from the Windows server to the Linux server using Macromedia's HomeSite the files seem to get corrupted. I've asked on some forums for HomeSite and people seem to think it should work fine with Linux servers. I'd rather not have to get the consultants back in if I don't have to.

Thanks in advance for your assistance in this matter.

Charlie, via email

There are slight differences in the way Unix and Linux systems save text files compared to Microsoft systems. Microsoft systems save line feeds as a Carriage Return character as well as a Line Feed, whilst Unix systems only use the Line Feed special character. There are many ways to convert these files and one such way is by using the *dos2unix* package, which may already be installed on your server. The syntax is as follows:

```
dos2unix -o file2 file1 file3
```

or

```
dos2unix -o *
```

There are further options to the *dos2unix* command for specifying another name for the original file to avoid overwriting the original.

I have also read unconfirmed reports that HomeSite is able to save

directly to Unix file format but I recommend checking HomeSite's extensive built-in help files.

Logwatch info

Q I use a standard Red Hat 8 install and I get a daily email with information from *logwatch*. Is it possible to get any further information out of this? Also, is it possible to adjust *portsentry*?

I'm more interested in getting the maximum information out of *logwatch*, adding logs to it etc.

Stuart, via email

A Yes you can get more information from *logwatch*. The main configuration file can be found at `/etc/log.d/logwatch.conf`. You can add almost any running service and set log levels from low to high.

Portsentry is also very configurable but is a little more complex – have a look at `/etc/portsentry/portsentry.conf`. Though *portsentry* can be set to be totally paranoid about others trying to connect to your server this is rarely a good thing to do. It may be best to leave it as it is unless you have a particular requirement.

Spam & FormMail

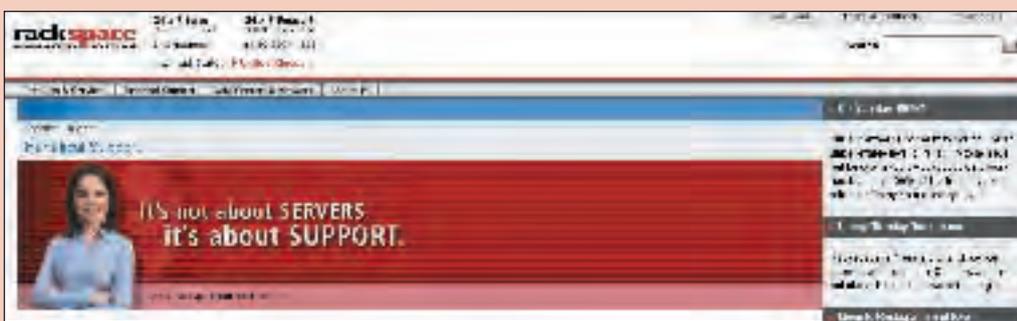
Q I run many virtual hosts on a Plesk-based server. Many of my customers use *FormMail.pl* to receive feedback from their users or customers. Recently I've seen a huge amount of spam being sent through my server and I've been told to check my *FormMail* scripts. I'm not sure what to look for though and with as

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 **Alan** – your new Zaurus PDA will be with you shortly!

Q I have a shared hosting reseller account, which I use to manage and host a few domains. It was easier to do it this way, as it gives me total control over my domains and it is quite cost effective (and I still have space to sell to cover costs).

My hosting provider is very good and has excellent support, but problems do occur, whether it's network problems in the data centre, hardware problems or DOS attacks on servers, it does happen that occasionally (thank goodness not often) my web sites are 'offline' or unavailable for browsing. My question is: is there anyway of automatically redirecting traffic to a (different) website where I can display a "Sorry, we are having technical problems, please come back later" sort of message/page, perhaps with a form so they can send us a mail, so we can email them as soon as things are sorted out. I would imagine that something could be done using the DNS settings to use a 'fallback' server/site, but I'm not sure if this can be done or how. Can you please explain if it is possible and how, or why not?

Alan, via email

A Although there are ways to do what you are asking, they are either difficult, imperfect or more than slightly expensive, which is why you'll mostly see this kind of functionality limited to big corporation sites.

DNS has no built-in functionality to check whether a service is up or down, it will pass an IP along for a name regardless. Here is an example:

If I want to visit your site and my ISP does not know your IP address, then it will look to your name servers and get the appropriate records. My ISP's name server will now cache this record. If anybody else who uses my ISP's name servers now visits your site they will use the cached entry. By default most domains on the Internet are cached for 24 hours.

It is possible however to set your Time To Live (amount of time the entry is cached), to only a few minutes. Then when your site goes down, you can change the DNS entries to point to the other server, and within a few minutes people should start hitting the new site with the holding page. The problem with this work-around is that certain DNS servers will override your TTL settings and cache the domains for longer. There's another snag – this setup also requires that your name servers are not at your current hosting provider and that you have full control of them (not all service providers are willing to reduce this TTL setting permanently).

A failsafe solution would be a globally load-balanced setup with servers in two datacentres and load-balancer hardware capable of doing global load balancing. This is probably way out of your budget though, and is normally only available at large dedicated hosting companies or co-location providers.

many customers as I have it will probably take a very long time.

Tim, via email

A Fortunately, it is fairly easy to find and replace any bad *FormMail* scripts on your server. At the time of writing it is safe to use version 1.92 of *FormMail* – anything lower I would seriously consider replacing. To find all the *FormMail* scripts you could simply use the locate command as follows:

```
[root@tinuviel root]# locate
FormMail.pl
```

```
/home/user1/www.user1domain.com
/cgi-bin/FormMail.pl
```

```
/home/user2/www.user2site.com/cgi-
bin/FormMail.pl
```

```
/home/user2/www.user2site.com/ht
docs/FormMail.pl
```

As you can tell in this example, there are three **FormMail.pl** scripts on this dummy server. To check the versions you could use the following command. You could even schedule this command in a *cron* job to mail you a daily or weekly report of all the *FormMail* versions.

```
[root@tinuviel root]# head -3
`locate FormMail.pl`
==> /home/user1/www.user1
domain.com/cgi-bin/FormMail.pl <==
#!/usr/bin/perl
#####
# FormMail Version 1.6
==> /home/user2/www.user2site.
com/cgi-bin/FormMail.pl <==
#!/usr/bin/perl
#####
# FormMail Version 1.6
==> /home/user2/www.user2site.
```

```
com/htdocs/FormMail.pl <==
#!/usr/bin/perl
#####
# FormMail Version 1.6
```

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!

« Flash for Linux?

Q I'm new to Linux but am enjoying using it so much (with the help of your mag!) that the only thing keeping me dual booting with Windows XP is Flash MX; for every other Windows app I've found an equivalent as good or, more often, better – but there's no equivalent even close to Flash that I'm aware of, even though the swf is an open standard (I think). So, after some initial difficulties, I have wine20030709 installed on Red Hat 9, wine-only, and am running Flash MX successfully – the only thing left to fix is system fonts which are tiny and unreadable. Wine does a very good job with fonts on simpler GUIs, but Flash is a veritable 747 cockpit of property panels, palettes etc and no fun to use at this stage. So I followed these steps which I got from the Wine documentation :

- Grab copies of the Windows bitmap fonts serife.fon, sserife.fon and vgasys.fon
- Extract with `font2bdf`
- Convert to pcf with `bdf2pcf`
- Stick them in your fonts path eg `/usr/lib/X11/fonts/misc/`
- Run `mkfontdir`, `xset fp rehash` and `/etc/init.d/xfs restart` to get the fontserver reset
- Edit the wine config file to remove the aliases "for the fonts you just created" which I presume are the ones that look like:

```
Default = "-adobe-helvetica-"
```

```
DefaultFixed = "fixed"
```

```
DefaultSerif = "-adobe-times-"
```

```
DefaultSansSerif = "-adobe-helvetica-"
```

which I duly commented out – no go. One thing I noticed is that all other files in that fonts dir are compressed and named in some strange format, eg `5x7-IS08859-14.pcf.gz`, so I gzipped the 'Windows' pcf fonts, with no more luck than you'd expect... Have you got any ideas about this? I need to save my eyesight!

James Burton, via email

A Compressing the .pcf files won't make a whole lot of difference. You'll want to check out the fonts using something along the lines of `xfontsel` so you can verify that the fonts you installed are actually available under X. Then you



Installing just the packages you need keeps your system buoyant.

can work on modifying the Wine configuration to reference those rather than the standard fonts. If the fonts in Wine look messed up, it may also be a good idea to look at some TrueType fonts and make those available to X via `xf86` so Wine can do some font smoothing and hopefully make things a little easier to read.

Network advice

Q I have been using Linux now for over a year and feeling somewhat adventurous I decided to bite the bullet and scrapped Windows 98 and installed Mandrake 9.0. Since that date I have tried Red Hat 8.0, Slackware 9.0 and SUSE 8.2. I am now sticking to SUSE 8.2. I have started various projects that I did not think I would ever do; eg CSS and website designing, Python programming, and of course, the inevitable networking!

And that is where my question comes in. I have read many HowTos and several articles on building a network but as of yet I am not networked! I can't get past the basics. What do I do first? Do I set up my firewall? I was going to use Smoothwall – is this a good idea. Do I set that up first? Then I was going to install SUSE 8.2 to a server but it doesn't come with Apache, MySQL or PHP as standard. Should I therefore get SUSE Enterprise Server? I have several PILs and I intend to install Slackware to them. When do I connect them in? I did get to the point of having two machines networked (or so I thought) but they wouldn't talk. The hub I have is a donated one. The lights come on so I assume it works, but how can I tell?

Richard Brown, via email

A The first step is to look at the NET-4 HOWTO. This looks at many elements of Linux network configuration, so should at least be able to get you to ping

between a few different boxes. Once you have your internal subnet configured, you can look at setting up a gateway, either using SUSE or Smoothwall, so you are able to access the Internet. If you get a link light on the hub, it at least means it's doing something, but without any information on the specific configuration on each of the network devices connected to the hub, it's rather difficult for us to determine if it is really working or not.

SUSE 8.2 does have Apache, MySQL and PHP available. You will just have to install the packages from the installation media. There is really no need for any Enterprise Server edition when you just need some basic web and database capabilities. During installation there should also be an option where you can select packages which you want to have installed, allowing all three sets of packages to be installed from the outset.

Linux multimedia

Q I'm looking for music player software, where it is possible to save the information (album title, songs titles, time) that we receive from [freeDB.org](http://freedb.org) or CDBB. I personally

have in my collection about 5000 CDs and my project is to build a database of all my CDs, and songs with songs time. If I can just listen to a CD in my computer and retrieve the information and save it, I will save a lot of time. I looked around and I didn't find any software in the Linux or Windows world, but I'm pretty new to Linux and maybe the information is saved somewhere on my computer when I'm listening to a CD and I don't know where. Maybe I can retrieve this information with MySQL, a spreadsheet or just a word processor (it's just data?) If you have any ideas or information, I'd be most grateful.

Sorry for my poor English, I'm French-Canadian and French is my first language!

Bernard Therriault, Rimouski, Quebec

A There are quite a few different Open Source options for you out there. One of the best we could find was *PHPmyCDMouse* (<http://sourceforge.net/projects/phpmycdmouse>), which runs under Apache with PHP support. You can also look at a variety of other options at <http://freshmeat.net/> [LXF](http://lxf)



GNU Sound is a multi-track audio editor and processor.



Xine is one of several movie and DVD players.

Essential disc info

Read this important information before you use your *Linux Format* coverdisc – CD or DVD. We've collated some helpful info to help you get the most from these jewels of data!

FINDING THE ESSENTIALS

Missing something?

As many of the programs on our discs are the very latest releases, they are often built on the very latest libraries and may depend on other packages your current Linux setup does not contain. We try to provide you with as many of these important supporting files and libraries as possible, though obviously we don't have space to include absolutely everything.

In many cases, the latest libraries and

other packages you might need will be included in the "essentials" folder on the disc, so if you are missing dependencies, this is the first place to look.

Package formats

Wherever possible, we try to include as many different types of package for an installation as possible, whether that be distribution specific RPMs, debs or whatever. Please bear in mind that we can only do this where space permits and when the packages are available.

We will, apart from exceptional or legally restricted situations, include the source files for any package, so that you can build it yourself.

Documentation

These pages provide helpful information on how to install and use some of the packages on the CD. Please note that many of the applications come with their own documentation, and there are additional notes and files in the relevant directories.

WHAT ARE ALL THESE FILES?

If you are new to Linux, you may find the profusion of different files and extensions confusing. As we try to give as many packages as possible for compatibility, there will often be two or three files in a directory covering different types of Linux, different architectures and usually source and binary versions – so which do you install? They can be identified by their filenames, and usually just by the file extensions.

Someap-1.0.1.i386.rpm – This is probably a binary rpm, designed to run on x86 systems.

Someap-1.0.1.i386.deb – The same, but a debian package.

Someap-1.0.1.tar.gz – This is usually source code.

Someap-1.0.1.tgz – Same as the above, tgz is abbreviated form of tar.gz

Someap-1.0.1.tar.bz2 – Same, but uses bzip2 compression instead of zip

Someap-1.0.1.src.rpm – This is also source code, but supplied as an rpm to make it easier to install

Someap-1.0.1.i386.RH7.RPM – A binary, x86 RPM designed specifically for Red Hat Linux

Someap-1.0.1.ppc.Suse7.rpm – A binary RPM designed specifically for SuSE7.x PPC Linux.

Someap-devel-1.0.1.i386.rpm – A development version.

INSTALLING FROM TARBALLS

A tar ball is a two stage archive. First the files are archived into a single file with tar and then compressed with Gzip or Bzip2. To unpack, cd to the directory you want to unpack it, usually your home directory and type one of the following two lines:

```
tar xzvf /mnt/cdrom/Desktop/progname/progname-2.1.0.tgz
```

```
tar xvf --bzip2 /mnt/cdrom/Desktop/progname/progname-2.1.0.tar.bz2
```

Use the first for Gzipped files, those ending in .tar.gz or .tgz, and the second for Bzipped files, ending in .tar.bz2 or .tbz2. Naturally, you change the paths to suit the location and name of the archive. and replace /mnt/cdrom with whatever is applicable to your system (eg /cdrom). This normally unpacks the archive into a directory of the same name, enter that directory with:

```
cd progname-2.1.0
```

To compile and install the software, type the following three commands:

```
./configure
```

```
make
```

```
su -c "make install"
```

The last line will prompt you for the root password, as this stage must be run as root. If you are already logged in as root, just type **make install**. This will give you a default installation. If you want to change any aspect of the install, type

./configure --help to see the options available. For example, you are usually able to change the default location with the PREFIX argument. When you have finished installing, you may remove the source files with:

```
cd ..
```

```
rm -fr progname-2.1.0
```

You should also log out as root, before you do anything you may later regret.

DEFECTIVE CDs

In the unlikely event of your 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**.

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. [LXF](#)

Coverdisc



Neil Bothwick is your guide through the wonders of this month's jam-packed *Linux Format DVD*. You'll be spoilt for choice among the 4GB of great programs!

The DVD also contains Gentoo Linux, but with many more packages. This makes the decision as to whether to copy the files from the disc to the distfiles directory more important, since there are over 2GB of them. If you have the space, it is definitely the easiest way. Otherwise the installation process is the same. Should you wish to install Gentoo on a computer without a DVD-ROM drive, it is easy to create an ISO image of the CD version, which you can then burn to CD-R. The following command will create the ISO image in the current directory.

```
sh /mnt/cdrom/Distros/Gentoo/mkiso
```

If you want to put it somewhere else, either CD to that directory first or give it as an argument.

While this is a reliable way of creating the ISO image, it does take a while as the program needs to scan the entire DVD to find the files it needs.

SYSTEM RPM-ANALYZER

Unless you use a distribution with advanced package management tools, RPM dependencies can still be a bit of a headache. This is especially true if you want to free up some space on your hard disk.

After a period of trying out packages, installing them with their dependencies, then uninstalling the package itself when you no longer want it, you can end up with a lot of those dependencies installed, even though they are no longer needed. Finding out which packages you can safely remove without breaking others can involve a long session of trial and error. What would help would be a way to visualise the relationships of the various packages on your system, so that you could quickly and easily see which could be removed without affecting anything else.



CD users only get KDE with Gentoo but you lucky DVD people get GNOME too, and a whole load of extra programs, more than enough to set up a system to do just about anything you can think of.

RPM-Analyzer does just that. The program displays information about the packages installed on your system in a variety of ways. Some of these are similar to those provided by other package managers, showing a list of packages with information on each one and a detailed description on the

selected package. Where this program can really help is with the tree views that show the relationship of packages. For any installed package, you can show either a list of the packages it requires, or a list of other installed packages that depend on it. The latter view can be particularly useful when performing housekeeping tasks, such as finding those no-longer-needed

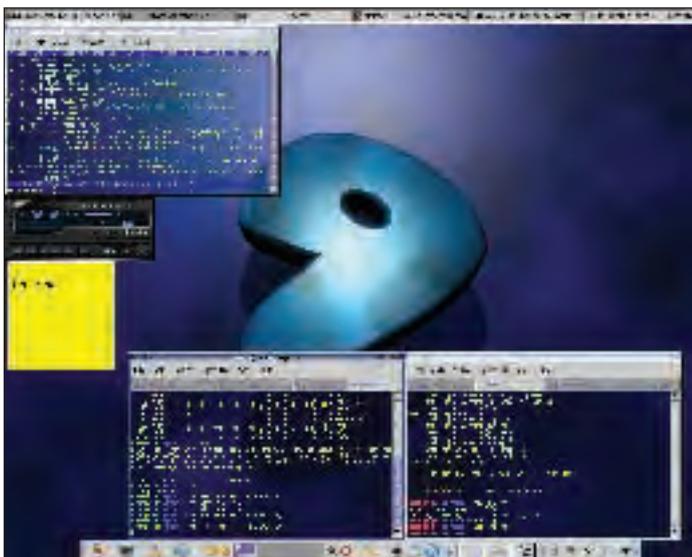
On the DVD



Wherever you see this logo it means there's related stuff on the DVD

IMPORTANT NOTICE

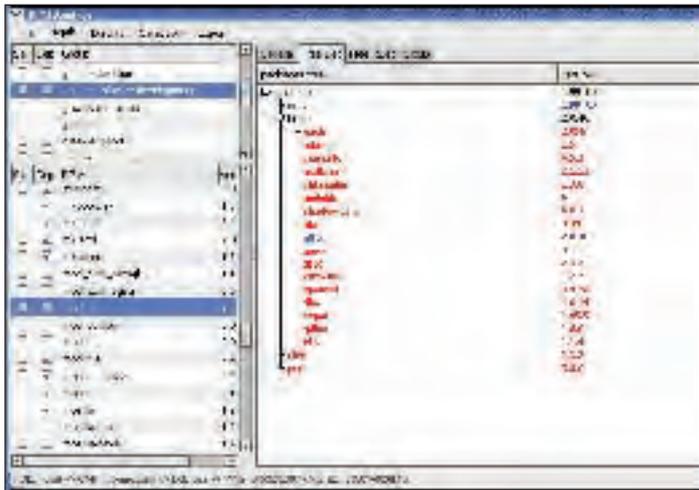
Before you even put the DVD in your drive, please make sure you read, understand and agree to the following: The *Linux Format DVD* is thoroughly tested for all known viruses, and is independently certified virus-free before duplication. We recommend that you always run a reliable and up-to-date virus-checker 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.



This is a Gentoo desktop running the XFCE window manager, one of many that are supplied on the DVD.

GENTOO FOR PPC

The source-based nature of Gentoo makes it less architecture dependent than most distributions. The *Distros/Gentoo* directory of the DVD contains a bootable ISO image for Gentoo PPC. I was not able to test this (if you'd prefer PPC software to be tested, just pop a G5 Mac in the post and I'll see what I can do!) but the PPC installation instructions are on the coverdisc and most of the source packages should be the same. It is possible that a few patches or packages specific to PPC are missing, so it would help to have an Internet connection while installing, and Portage will download these as it needs them.



Keep track of which RPM packages you need and which are simply a waste of disk space, with *RPM Analyzer*.

obscure libraries that you installed two years ago while you were trying out some new software that didn't do quite what was wanted.

RPM-Analyzer is written for and tested on Red Hat, but it should work with most RPM distributions. If installing from the **noarch** RPM doesn't work for you, try rebuilding from the source RPM file.

DISTROS RUNT

USB flash drives are becoming ever cheaper and larger. They usually take the form of a key fob, so you can carry your MP3/Ogg Vorbis collection around in your pocket, although for real geek-appeal there are even watches with built-in USB storage – get yours from www.thinkgeek.com/gadgets/watches/5eec/. Whatever format they take, these drives are now large enough to be used for more than transporting a few files. They now have

the capacity to install a very reasonable Linux distribution. RUNT (ResNet USB Network Tester) is such a distribution. It is based on Slackware and is similar to Ziplack.

There are two files included on the DVD: a zip file to be extracted to a 128MB (or larger) USB drive, and a bootable floppy image. The floppy disk is needed the first time you use RUNT, but if your computer supports booting from USB devices, you can run *LILO* and boot direct from the flash disk subsequently. If your computer will not boot from USB you will still need the floppy disk. This makes RUNT a little less convenient than carrying something like Morphix-Light on a 3" MiniCD, but you have the advantage of a writeable disk.

The USB drive needs to be formatted with a FAT filesystem, which is how they are normally supplied. The disk doesn't need to be empty, but it does need around 115MB of free

I DON'T WANT GENTOO!

But there's still lots for you here...

Even if you have no interest in Gentoo, the cover discs still provide a lot for you. Because Gentoo uses the original source tarballs as the basis for installing software. These can also be used with any other Linux system, just by unpacking the tarball and

following the instructions in the **README** and **INSTALL** files. So even if you don't want to try Gentoo yourself, you'll still find a CD full, or around 2GB if you have the DVD edition, of Linux software that should work with your current distribution.

space. Mount the disk, **cd** to its mount point and type

```
unzip /mnt/cdrom/Distros/RUNT/runt-2.0.zip
```

to install the files. Then put a floppy disk in the drive and type

```
dd -of=/dev/fd0 if=/mnt/cdrom/Distros/RUNT/runtboot2.img
```

to create the boot floppy. Now unmount the USB drive and reboot with the floppy still in the drive.

DESKTOP LINUX EASY ACCESS KEYBOARD

Many keyboards have a number of extra keys, often referred to by hardware manufacturers as 'Internet' or 'Multimedia' keys. In reality, they are just extra keys with fancy logos and labels, their functions depend on the software used to interpret the responses from those keys. The keyboards usually come supplied with Windows software to map these keys to specific functions, and here is the equivalent for X.

There are three programs here. The most important is *lineakd*, a daemon that watches for presses of the extra keys and executes whatever commands you have told it to run. This is all set up in the configuration

file, which is where the other two programs come in. *Lineakconfig* is a GTK program to configure *lineakd* – it's where you tell the daemon which keyboard you have and what to do when each key is pressed.

Klineakconfig performs the same function, but using a KDE/QT interface. Later versions of this software will only have the KDE interface for configuration, although the program itself works equally well whatever window manager you use. The programs are supplied as source code, Debian packages and RPMs, use whichever installation method suits you. The RPMs were created on Mandrake, they should work with other RPM distributions. If they do not, there are source RPMs with which you can rebuild binary RPMs to suit your system.

SYSTEM VCool

Athlon processors are known for their high operating temperatures. This is mainly a consequence of their high power consumption. The power consumption increases running costs, while the heat it generates requires extra cooling, which usually results in more noise. There is also a question over how much this affects reliability, not only of the CPU, but also of all the other components that have to operate in a hotter environment. What is less well-known is that Athlon CPUs have a power-saving mode that reduces power consumption and heat output when the processor is idling, but this is normally disabled in the motherboard's BIOS. *VCool* is a small program that enables power-saving on VIA, nVIDIA, SiS and AMD motherboard chipsets. It is possible for this setting to make a system unstable, which is why motherboards have it turned off by default, but for many systems it is stable, as well as making things much cooler, so why not give it a try? [LXF](#)

WHAT'S INCLUDED WITH GENTOO?

Over 2000 packages and patches!

The distfiles directory on the CD contains 341 packages and patches, the DVD has over 2200. The CD has a base Linux installation, with a choice of kernels, XFree86, KDE, ALSA, CUPS, Java and anything else you need to set up a fully functional desktop system. It also includes several major applications, such as *Mozilla*, *The GIMP*, *Apache* and *MySQL*.

The DVD has far too many packages to list, so here is a selection of some of the highlights. There is much more than this on the disc, and there are also all the libraries and other packages that many of the programs require.

XFree86
KDE
GNOME
NVIDIA drivers
ATI drivers

Scribus
OpenOffice.org
KOffice
Abiword
Gnumeric
Gnucash
Mozilla
Galeon

Evolution
Downloader for X
Sylpheed
Gaim

Mplayer
Ogle
Xine
XMMS
The GIMP
Gphoto

GCC
Python

Perl
Ruby

MySQL
Postfix
ProFTPD
Vsftpd
Zope

WINE
X-CD-Roast
K3b
and many, many more...

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

LUGS OF THE MONTH

Menai Linux Users Group/Menai Grwp Defnyddwyr Linux

Croeso i Menai LUG. Mae Menai LUG (Grwp Defnyddwyr Linux) yn agor i bawb yn ardal y Fenai sydd efo diddordeb mewn darganfod rhagor am Linux ac meddalwedd rhydd. Cewch ymuno a'r rhestr ebost drwy fynd yma: <http://mailman.lug.org.uk/mailman/listinfo/menai>

Ydych chi eisiau cael Cymraeg ar sgrin y cyfrifiadur? Mae yna gyfle! Mi allwch chi helpu gyda cyfieithu penbwrdd Linux i Gymraeg trwy fynd i'r gwefan www.kyfieithu.co.uk i

www.gwelywiwr.org. lle gewch chi fwy o fanylion am y menter. Os ewch chi i "Newyddion" mae yna erthygl am y fenter. Oherwydd gwyliau y haf ni fydd gyfarfodydd Menai LUG... Ond mi fydd y cyfarfodydd yn ail-ddechrau yn mis Medi!

Welcome to Menai LUG. The Menai Linux Users Group is open to all in the Menai area who have an interest in discovering more about Linux and free software. You can join the mailing list by going here: <http://mailman.lug.org.uk/mailman/listinfo/menai>

Do you want to see Welsh on the computer screen? There is a chance! You can help with translating a desktop into Welsh at www.kyfieithu.co.uk and www.gwelywiwr.org where you will also find more info about this. If you go to the 'News' section there is an article about the translating. Because of the Summer holidays there'll be no Menai LUG meetings... But the meetings will restart in September!

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.luvasn.au
Contact luv-committee@luvasn.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.aalslug.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 afyde@balug.org
PAKISTAN
URL www.linuxpakistan.net
Email tux@clug.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

Free Software is obviously of importance to not-for-profit organisations. Jono Bacon needs to discuss this further.

Last month we started looking at how charities can make use of Linux. Charities seem a natural target for free software due to the fact that they need to preserve every penny for meaningful uses. This month we will identify some of the issues we can explore when advocating Linux to charities.

The first thing we need to understand is that although charities need to save money, they also need as a solid a system as anyone else. Charities are not willing to skimp on functionality just to save some money. In terms of software usage, it is likely that we have it covered. Productivity can be achieved with *OpenOffice.org*, *The GIMP*, *Scribus*, *Mozilla* etc, development with PHP, Perl, C/C++, servers with *Apache*, MySQL, PostgreSQL etc. however, one of the challenges to charities is support.

Support is an issue where we as a community can really help charities. It is likely that a charity will only need some help with new topics (such as when they need to install software/hardware). This is also an area where LUGs can really come

into their own. Remember, as a community we need to work together!

To help facilitate the discussion of assisting charities in using Linux, I have set up several resources on my website. I could not find any suitable info to give to a charity to explain what Linux is and how it can help them so I decided to write one myself and developed the 10-page PDF at www.jonobacon.org/writing/research/linuxforcharities.pdf – feel free to use it in your advocacy! Part of the problem is finding charities that use free software already, and to remedy this I set up a page to allow charities to show they use free software: www.jonobacon.org/projects/charitiesregister/. Finally, as there's only limited space in the magazine to discuss advocacy to charities, I have put up a forum at www.jonobacon.org/forums/, and I would be interested to hear your thoughts. This is definitely an area in which free software can make maximum impact. Next month we are going to focus on how Linux and free software can help local businesses.

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

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

Contributors

Hans Huberland, Rachel Borg

Photography

Powerstock, James Wilson

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

Welcome

A real-world Linux supplement for IT professionals and IT managers

In *Linux Pro* this month, we concentrate on three areas of highest concern for Expo attendees – Linux security, Linux training and Linux storage. Rackspace sysadmin Hans Huberland examines official Red Hat certification; 7Safe shows Paul Hudson how to think like a hacker to plug network vulnerabilities, and Xinit gives you an insight into their operation and why Blade technology is the future.

Linux Format and *Linux Pro* are proud to be Premier media partners of Linux Expo UK 2003, and this special edition of *Linux Pro* also doubles as the showguide for the event – just turn it upside down and read from the back. I will be speaking as part of The Great Linux Debate myself – if you've got an opinion about Linux that you would like to hear discussed in an open forum, come along and contribute.

We're on **Stand 145** – please stop by and let us know what you think of both *Linux Format* and *Linux Pro*, and give us suggestions for things you would like to see covered in future issues – we always relish the opportunity to get to know our readers better! Whether you just discovered Linux last week, or administrate a network of thousands of users, staff from the mags will be on hand for you to poke fun at, swap helpful Linux advice or just chat about all things Linux, Free Software and Open Source-related (they tend to be more talkative when plied with beer!) Enjoy the Expo – we'll see you there!

Nick Veitch Editor

nick.veitch@futurenet.co.uk

CONTENTS

RHCE CERTIFICATION

Distro-specific Linux training **p4**



BEAT THE HACKERS!

7Safe shares its experience **p6**

XINIT CLOSE-UP

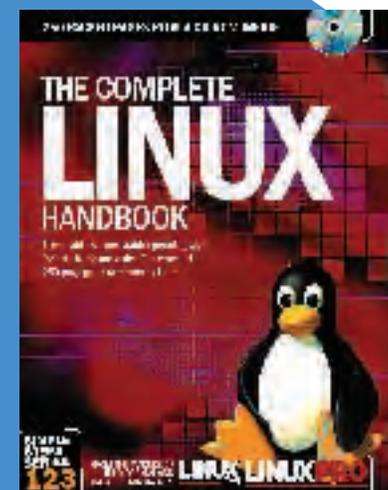
Linux across the board **p10**

LINUX

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



Subscribe at
LINUX EXPO
2003 and get
'The Complete
Linux
Handbook'
FREE!



WORTH £14.99!

Offer based on 13-issue subscription @ £54.90 [CD] and £59.90 [DVD]. Offer UK only.

**SUBSCRIBE
& SAVE £24.47**
[*savings based on DVD edition]

VISIT US AT THE LINUX FORMAT STAND 145



Red Hat Rapid Track RHCE Training and Certification

Vendor-specific training has not had a particularly good reputation in recent years. This can probably be largely attributed to the fact that it's near impossible to teach the intricacies of a complex system like an OS in a number of days or even several weeks. Red Hat is using the RHCE certification as a means to, "Prove what you can do," as opposed to, "Prove how much you can cram into your brain before exam day."

This is a Rapid Track course and is not targeted at newcomers to Linux system administration, but instead at sysadmins who have a solid understanding of the OS, and who would like to round their knowledge off with aspects they're not exposed to on a regular basis. Red Hat offers several other courses aimed at differing experience levels, starting with *Linux Essentials* through to *Kernel Internals* and *Device Driver Programming*. Red Hat is putting a lot of emphasis on its claim that this certification is different and that an administrator capable of attaining it is sufficiently qualified for the task at hand. This is backed up with a long and very practical test with a pass mark of 80%, and an unofficial pass rate of about 50% for first-time candidates.

Course Contents

The syllabus covers vast tracts of system administration in only four days so you should expect that not much time can be spent on any one topic. The range of topics covered is broad but there are a few arguably important areas that are left untouched. A lot of emphasis is placed in NIS and NFS despite their generally non-secure status. Whilst mentioning that LDAP is a far more powerful and secure service than NIS it is not covered at all. Surprisingly, kernel compilation is

Sysadmin HANS HUBERLAND goes back to school with Red Hat and its much-talked-about Rapid Track RHCE training and certification.

RESIDENTIAL TRAINING COURSE

Training and certification in Red Hat available for several different levels.

PROVIDER Red Hat
LOCATION International
DURATION 4 days + 1 day exam
COST £1600

only briefly mentioned and not covered at all in the labs. Despite these flaws, overall the training goes into enough detail on each topic to allow a moderately experienced admin to figure the rest out. Needless to say, enough info is covered to take a student through the exam, but the candidates will certainly need to add to this with specific research for their own particular real-world implementation.

Materials and Instructor

Red Hat's premier UK training site is at its EMEA HQ in Guildford, Surrey. The classrooms are surprisingly well-equipped, with a sufficiently powerful Dell system for each student, as well as an effective KVM system by means of which the instructor can grab control of any student's system to lend a hand or mirror his display onto every student's monitor for demonstrations. Up to 12 students per class allows a level of personal attention from the instructors.

The instructors are all Red Hat employees, so they have a certain amount of insight into what the company's plans are. This is not to say that you will be privy to confidential info, but they are able to put into plain English what already appears in marketing speak on their website. From my experience (your mileage may vary) the instructor's understanding of the topics at hand was thorough, and at times when they do not know the answer, they are happy to admit it and do some research.

Red Hat's trainers have good communication skills and established a strong rapport with the students. Despite the quantity of material covered, the instructors managed the time extremely well, ensuring that there was sufficient time on the last day to cover example practical questions without having to skim over other sections and play catch-up.

Exam

All RHCE exam participants are required to sign a binding non-disclosure agreement prohibiting them from discussing exam questions and practical scenarios. This taken into account, Linux Pro has been permitted to describe the exam process, so no telling on us! The exam consists of three sections: a two-and-a-half hour troubleshooting lab, a one-hour multiple-choice test and a two-and-a-half hour installation lab. An average score of 80% over these three sections is required, and a result of less than 50% on either of the first two sections results in failure; a score of at least 70% is required for the installation section.

The troubleshooting lab presents students with their machines in a non-booting state. Using only the rescue disks provided, students have one hour to complete the first four of eight problems. These first four are relatively straightforward, and all four must be resolved successfully to pass the exam.

The multiple-choice section is a fairly self-explanatory: 50 questions, four answers to choose from, one correct answer. This is the only section where they tend to go into the precise technical details. The practical questions can be solved however suits you (provided you meet the objectives), the multi-choice requires that students pay attention in class and spend time after the classes reading over the notes provided.

The installation section is debatably the most difficult, not by virtue of the content but by sheer time restrictions. Students are handed a 4-page list of requirements and let loose. All the requirements are covered in the 4-day course and there are no nasty surprises, this section tests how comfortable students are with the individual aspects of the course. Candidates are free to install any documentation but having to refer to this material more than just once or twice briefly could result in running out of time. This section is also divided into two sections, a local services section and a



Red Hat's training lab is well-equipped and networked to the tutor's machine to avoid the sort of 'over-the-shoulder instruction' that characterises many IT training courses..

network services section. Candidates who pass the local services section but fail the network services section can still acquire the lesser RHCT (Red Hat Certified Technician) title.

Conclusion

This exam deserves respect. It puts candidates under the same pressure they would be under in a real-world environment – man pages are available, but phoning a friend is A Bad Thing! The course itself compares fairly well compared to the other vendor-specific ones, but still suffers from a few of the drawbacks that would be expected to be associated with such focus on a particular product. The exam is well worth it, although initials behind somebody's name don't guarantee their skill, anyone who has passed this exam will at the very least have the initiative to solve problems and can work well under pressure. ■

GENERIC ALTERNATIVES

Other Linux certification standards

THERE ARE VENDOR-NEUTRAL LINUX training courses out there. The two most famous examples are probably Linux Professional Institute's LPIC and CompTIA's Linux+. There are both advantages in being distro-agnostic. For instance, if your training is more general, it will be of use across many different Linux systems; however, it is unlikely that you would get the sort of in-depth knowledge about likely problems experienced on particular system distro that more focused training would provide. Both of the alternatives below have something important in common with RHCE: all three emphasise a combination of Linux knowledge and the correct application of that knowledge.

LPI Certification

www.lpi.org/

This consists of three levels: Junior, Intermediate and Advanced. The junior level

LPIC1 focuses on local configuration including X, kernel module setup and CLI navigation. LPI Level 1 is wholly encapsulated in the O'Reilly book, *LPI Linux Certification In A Nutshell*, (ISBN 1-5659-2748-6) which is handy if you are a good self-learner. The Intermediate level is largely on par with the RHCE, it focuses on network server configuration for medium-sized networks. The advanced LPIC3 is still under development but will focus on managing large corporate environments and busy, mission-critical Internet servers. For more on LPI certification (available in the UK through Training Camp), see *LXF43's Linux Pro* and visit the .ORG village at Linux Expo 2003.

CompTIA Linux+ course

www.comptia.org/certification/Linux/default.asp

A new addition to the Computer Technology Industry Association's training portfolio.

The CompTIA Linux+ course tests seven domain areas in the following proportions:

- Planning and Implementation (4%)
- Installation (12%)
- Configuration (15%)
- Administration (18%)
- System Maintenance (14%)
- Troubleshooting (18%)
- Identify, Install and Maintain System Hardware (19%)

Some reports on the Internet indicate that some people who have taken it think that the exam is relatively easy, and that even people who are not very confident using Linux would have little problem passing. It would of course be inappropriate for *Linux Pro* to comment on this further, as we have not examined the course in depth, but we intend to do so in the future. The course's intent is stated as being to "validate the knowledge and abilities of individuals with at least six months of practical Linux experience".

Think like a hacker to

BEAT THE HACKERS

The world has never seen a greater need for secure computing, and yet nearly every day another company gets hacked because it didn't properly secure its network. **PAUL HUDSON** finds out about training that can combat this...

Cyberterrorism is more than just a buzzword used by the computer industry to scare you into parting with your money. For hundreds – maybe even thousands – of malicious hackers around the world, it's a way of life that drives them to pit their computer skills against network administrators of high-tech companies who believe themselves to be safe behind their array of firewalls.

While it's true that, historically, Unix-based systems have been a great deal more secure than their Microsoft equivalents, all too many people who move from Windows to Linux believe that they are getting a 'set-and-forget' network – install Linux, add in a few servers such as *Apache*, and you're safe. Not quite...

Zero Cool

The truth is quite simple, as Dennis Hughes of the FBI succinctly put it: "the only secure computer is one that's unplugged, locked in a safe, and buried 20 feet under the ground in a secret location, and I'm not even too sure about that one". While that might make the whole concept of computer 'security' seem like a farce, the situation is often a great deal worse.

Computers, generally speaking, get less secure as you install more software on them, and modern distributions such as Red Hat and Mandrake have minimum footprints approaching a gigabyte. Given that you can get a working Linux system with a variety of services up and running in under 100MB, there's quite clearly a lot of program code sitting around that is just waiting to be exploited.

While it's true that it's almost impossible to be absolutely secure, that doesn't mean it's not worth trying. The computer security industry is worth billions worldwide – and, unlike many other parts of the IT business – has shown little sign of slowing its explosive growth. However, simply purchasing and hooking up a pile of shiny black

boxes with flashing lights gives you no guarantee of security – at the end of the day, the security of your network is only as good as its weakest link.



Turning the tide

The most efficient and secure way to keep your network as safe as possible has always been to hire a top-notch security professional who understands the risks and is able to design and implement a strategy that keeps hackers out and data in. Of course, it's not possible for every company to hire someone like Niels Ferguson or Bruce Schneier, but, fortunately, there is an alternative to hand.

7Safe, a leading light in security training in the UK, offers a comprehensive security training course designed to teach students the hacker's mind set, and also how to use the tools that hackers themselves use. Through a variety of practical exercises and hands-on labs led by qualified instructors, the goal of the course is to educate students in the art of ethical hacking so that they can properly defend their own systems against attack.

The course examined in this article is *Hacking Insight: Hands-On*. Over the length of the three-day course, usually set in laid-back surroundings where the students are staying just a short walk from their classroom, a variety of topics are covered, including the typical techniques used to gain or force access into a system, how hackers conceal their tracks, how to bypass firewalls, and how to spoof TCP/IP packets. As well as this individually tailored training aimed at systems professionals with the day-to-day responsibility of keeping a network secure, the other standard courses offered by 7Safe are:

- *Hacking Insight for Managers* – a one-day non-technical course for decision makers
- *Building, Breaking & Securing Wireless Networks: Hands On* – a one-day practical course where delegates set up WiFi networks and break into them using the tools of a war driver-cum-hacker, then secure them.

HACKERS AND CRACKERS

HISTORICALLY, THE difference between 'hacker' and 'cracker' was that crackers were malicious whereas hackers were not. Sadly, thanks to Hollywood and much ill-informed sensationalism by television news and the tabloid press, this distinction has now largely disappeared from the public lexicon, and the term cracker is rarely used.

Rise of the machines

Perhaps the most important element of the 7Safe course is that it is largely hands-on – a number of hackable Linux and Windows machines are available so that students are able to work by themselves or in small teams to find and exploit real system glitches. For example, one exercise students need to complete has them using *netcat* to mimic how hackers launder connections to hide their identity. These exercises aren't just paper tests by any means – if students haven't learnt their stuff, they can't just pick from a list of options.

By breaking down tasks into easily grasped units, and encouraging students to discuss problems, solutions, and ideas at any time, the end result is that the course flows smoothly from the beginning to end, with many students opting to take the final exam confident in their studies.

Day one explores the history of hacking, teaching students the mindsets and motivations that inspire hackers to create chaos. The first technical topic



VIEW FROM THE TRENCHES

From theory to practice – how much difference does the course actually make?

TO GET THE VIEW FROM PEOPLE WHO HAVE actually sat the course, we spoke to Stephen Coombs, the Computer Systems and Security Administrator for the Defence Communications Services Agency (DCSA)

LINUX FORMAT: What prompted you to take the course?

STEPHEN COOMBS: System security has to be a major issue for us as part of the Ministry of Defence. Having taken the ISEB Principles of Information Security Management through Academy Training this course was an excellent and natural follow-on to gain a more practical insight into the methods used by hackers to gain access to systems. If you know how it's done, you have a better chance of preventing it.

LXF: How did you find the course?

SC: Fascinating. It takes a real in-depth understanding of operating systems to be a hacker. Under normal circumstances and without

training like this we would never come across such esoteria as root kits. There are some very clever people out there and all sysadmins can really ever do is play catch-up.

LXF: How has taking the course affected your work practices?

SC: Remember the two principles of computer security:

- 1 Don't buy a computer.
 - 2 If you must buy one, don't turn it on.
- We could provide a 100% secure computer system, but to do so we'd have to lock it down so tight nobody could access it and it would be completely useless. Alternatively, we could provide one that anyone could access however they wanted to. That would be equally useless as we couldn't guarantee data integrity.

In a practical sense, this course has helped us to develop the fine line between the two extremes that all System Administrators have to walk to provide our customers with a useable but secure system. We have a better understanding of how much free rein we can allow users without possibly compromising data integrity – and where the weak points are that we need to monitor. I can't be more specific than that because of the classified nature of the data held on our systems.

LXF: How would you rate the quality of delivery, course materials and professionalism of the course?

SC: Highly. Six of us attended the on-site course, all of us enjoyed it and all made the same criticism: we wanted to stay for a fortnight instead of three days! I can't recommend it enough.

“Without training like this we would never come across such esoteria as root kits...”

COVER FEATURE HACKING PREVENTION

covered, inline with what hackers would learn, is the technique of probing potential victims, with particular emphasis given to how hackers use the power of deduction and inference.

Day two moves into the area of vulnerability assessment, first from the perspective of a systems administrator, and then from a hacker's perspective. It is here students will learn how to perform and interpret penetration-testing reports in order that they can properly understand what vulnerabilities exist. Once the students understand where their potential vulnerabilities exist, the teaching of countermeasures begins. This topic, crucial to the course, is a combination of methodology, research, knowledge, and tools in order that the students are able to easily look for telltale signs of system compromise. At this point, the emphasis is very firmly on system defence – the students are taught to find their vulnerable points, and how to work around them.

The last day, day three, advances onto more complicated techniques and the use of other tools that further the goal of a hacker: The students learn how trojans are made and their payload executed invisibly, how to use root kits to conceal their presence on a system, and also how firewall and intrusion detection systems are evaded.

Judgement day

At the end of the course, the students are offered the chance to sit a final, hour-long exam that covers around 50 questions taken from the course content. Successful students are awarded the Certified Security Testing Associate (CSTA) certification, which allows them to progress towards the Certified Security Testing Professional (CSTP) assessment at a later date if they choose to.

Many companies in the security business sell their wares by scaring customers. It wasn't such a long time ago when an advert for a Commodore Amiga virus killer used a large tombstone – pretty much the epitome of the hard sell! What 7Safe does differently is to provide a program designed to calm fears of security by equipping people with the training, understanding, experience, and materials they need to develop and more importantly /sustain/ a corporate security policy that works.

As more and more companies realise that the security of their information and network is absolutely critical to their survival, it's possible that the grip hackers have on the Internet may start to loosen. Until then, 7Safe's team of experts stand ready to give you the knowledge you need. ■

VIEW FROM THE TOP

Dan Haagman, a technical trainer at 7Safe, answered our questions...

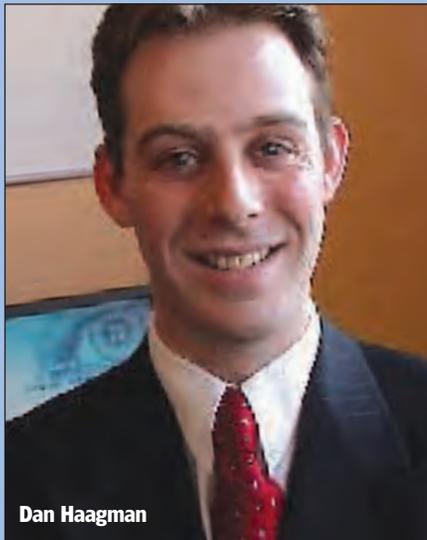
LINUX FORMAT: WHAT ARE 7SAFE'S MAIN goals with its security courses?

DAN HAAGMAN: 7Safe's courses are designed to show delegates just how real the threat of hacking is and how to defend an organisation and its assets. Our objective is to 'demystify' hacking by having them actually do it over three days in our class lab. By understanding what a hacker does (or what they can potentially do) enables proper layers of defence, controls and monitoring to be put in place and managed.

One of our main objectives is to impart our specialist knowledge with those responsible for maintaining and managing security. Hacking is seen as a bit of a dark art, and therefore requires understanding and a degree of expertise. *Hacking Insight: Hands-On* aims to build that expertise and expose the secrets of hacking for the purposes of effective defence. The majority of the tools featured in this course are widely available on the Internet, as the hacking community is generally willing to share their secrets with others.

LXF: What level of knowledge is required for people to take the course?

DH: Our course requires a basic understanding of TCP/IP with a background in Windows and/or Unix. No two courses are the same, as we tailor each course to suit the needs of the group at any one stage and we often find ourselves exploring specific student scenarios. Course materials are written in such a way as to accommodate those



Dan Haagman

who are relatively new to security concerns whilst providing a challenge to those with more experience and expertise.

LXF: How much detail does the course go into?

DH: The course is designed to be very practical and a hands-on learning experience. As such, it is very detailed, covering a significant number of tools and exploits. More importantly though, delegates learn to appreciate how a hacker draws together information to coordinate an attack and how combinations of tools are often used.

Hours of teaching are from 09:00 to 17:30 for all three days. We strongly encourage open discussion of the material and the sharing of

experience, which puts much of the course content into perspective. The course is intensive and fast-paced with 40 or so practical exercises used to consolidate the teaching material. We often find that delegates stay behind to ask further questions or run specific scenarios by our instructors on a one-to-one basis.

Following course completion, we also give out our contact details in order for those with further questions to receive some support and help in the future. A good majority of delegates do get in touch to ask about certain tools or for advice in consolidating the material in their labs, which we encourage everyone to build.

LXF: How much hands-on practice are delegates asked to do?

DH: *Hands-On* is designed to be just that. Hacking requires skill as does defence against hacking. By learning the methodology and applying it in a security-testing scenario, you can enhance your ability to secure a network or resource. We therefore spend time discussing a tool or technique and then move straight to using it in the classroom lab.

The machines are networked in pods whereby each student learns to 'hack' into their partner's machines and *vice versa*. Each exercise is thoroughly documented and explains step-by-step what is being done and achieved. Our trainers also work with each student pair to reinforce the material and draw out additional points for each scenario.

CUT YOUR TEETH

So, you think you know Linux security?

TRY YOUR HAND AT THESE EXAMPLE questions from the 7Safe exam...**1** In modern implementations of Unix, encrypted passwords are held in the /etc/shadow file

- A True
- B False

2 When port scanning using *nmap*, which type of scan is the most likely to be recorded by the scanned machine?

- A TCP ACK scan
- B TCP SYN scan
- C TCP connect scan
- D UDP scan

3 TCP wrappers can be used on Unix Systems. Which of the following comments is not true

regarding TCP wrappers?

- A They increase authentication beyond merely user name & password
- B They can be used to prevent spoofing of host names
- C They make use of the /etc/hosts.allow and /etc/hosts.deny files
- D They can't protect against TCP sequence number guessing

4 When UDP (User Datagram Protocol) port scanning a Linux machine, a hacker may not receive a response because:

- A Linux does not support UDP by default
- B A penalty is imposed by the kernel if a UDP port scan is detected
- C The Linux kernel has implemented the RFC describing UDP scans incorrectly.
- D UDP is connection-oriented

If you found yourself scratching your head, it's clear your security skills need brushing up with 7Safe. Luckily, as a *Linux Pro* reader you can claim a 10% discount on all 7Safe courses booked before 31 December 2003 and quote *Linux Format* magazine – that's £1215+VAT for the three-day *Hacking Insight: Hands-On* course (normally £1350+VAT), £495 for the one-day *Hacking Insight for Managers* course (normally £550+VAT), and £495 also for the one-day *Building, Breaking, and Securing Wireless Networks: Hands-On* course (normally £550).

All courses include lunch and refreshments, course booklets and exercise manuals where appropriate. For more information, visit the 7Safe website, www.7safe.com, where you can download the brochure as a PDF; or send an email to education@7safe.com for more info – and the answers to these questions!

LXF: What's the success rate of the course?

DH: It is often the case that delegates take away the concepts learned and immediately start applying them with the aim of protecting their networks. Our aim is to show the implications of being hacked and to raise awareness of the countermeasures that can be put in place. There has been a recent surge in enquiries for *Hacking Insight: Hands-On* as people are becoming aware of the impact of the latest spate of Internet worms, as are their managers, who have experienced measurable financial loss as a result. Delegates come from all types of organisations, including both the public and private sector.

LXF: How much interest is Linux generating?

DH: There appears to be a great deal of interest in Linux-based security. It is an easy assumption to make that hacking is focused mainly at Microsoft products. Security should be carefully considered for any operating system.

The various Linux vendors are certainly tightening up security at a very simple level by shipping default installations with numerous services switched off. Sadly, it is still not a case of "set and forget" as security flaws and exploits are constantly appearing in ALL operating systems. Misconfigurations are also another source of vulnerability that we believe are set to grow as software becomes more complex and the demands placed on Linux admins increases. Much of the course is based around using Linux as

a platform for many of the hacking tools. Linux is an extremely robust, flexible and highly configurable OS. Because of this, portability of hacking tools written in languages such as C are used across a wide range of Linux flavours such as Red Hat and SuSE. We use a couple of versions of Red Hat in the course for flexibility and to demonstrate where security flaws could exist. Linux is not only used as the platform for launching the attacks, but also for being attacked.

LXF: What kind of learning materials do you give out on the course?

DH: All slides used in the course are issued to delegates and these are designed to be perfectly usable as an *aide-memoire* later on.

Equally important is the exercise book that each delegate gets to take away. It forms a comprehensive resource for each exercise in the course and will read as a stand-alone document with explanation, rationale and pointers. All the commands and switches used with the hacking and security tools in the course are documented, making it an invaluable resource for use back at the office. References for each tool are also given at each and every point for hacking tools used.

Finally, throughout the course we make use of 'props' to demonstrate and reinforce the learning objectives. For example, as well as kernel-level keystroke loggers, a hardware keystroke logger called *KeyKatcher* is available for delegates to experiment with.



Attendees get to keep slides from the course for their own reference.

LXF: What examinations are done, and what's the pass rate?

DH: At the end of the course there is an optional one-hour, multi-choice examination. It is computer-based, giving the result immediately upon completion together with the student's score. Successful delegates will be awarded the CSTA (Certified Security Testing Associate) certification and this will qualify delegates for entry towards the CSTP (Certified Security Testing Professional) assessment. The latter is a half-day practical assessment of a candidate's ability to perform penetration testing in a simulated environment. Completion of the course also counts towards maintenance of ISC2's CISSP and SSCP qualifications, attracting 24 CPE (Continuing Professional Education) credits.

Figures to date show that 65% of delegates obtain or exceed the pass mark of 70%. The highest score has been 87% and it generally shows that those who revise end up passing the exam.

Xinit close-up

They use Linux across the board, donate hardware and code to various open source projects, and even employ a Debian developer. **PAUL HUDSON** takes a look at the future of storage systems and servers through the eyes of Xinit...



Despite reviewing Xinit's Sharq SPS234 about six months ago, it still has the highest benchmark results of any machine we've reviewed, and is in fact the only server we've reviewed to date that scored 10 out of 10. Of course, technology has moved on, and the SPS234 model we reviewed is no longer top of the Xinit range – Xinit has been refining and upgrading its server to include newer hardware.

At the same time, Xinit is more than just a server manufacturer. Back in the *LXF42* review, we mentioned that Xinit offers a choice of either Red Hat, SuSE, or Debian for its machine, which is something yet to be matched by many other companies. The reason the company is so confident it can offer solutions on all these platforms is because it is one of the few manufacturers to commit strongly to Linux across the board, which puts Xinit in a somewhat enviable position – not only does it take the extra time to design and test the systems with all the varieties of Linux they have at their disposal, but Xinit are also able to extend its technical support to include the distributions they send out.

The latest systems to appear over the horizon from Xinit involves storage servers, in the form of its SharqStor series of machines. Through the planned use of Serial ATA, Xinit is hoping to make available machines with up to 6TB of high-speed data capacity. The whole system is to be driven by a custom-written storage management solution based upon various open source components, and we hope to be able to bring you an exclusive review of this next issue.

In the meantime, in order to find out more about Xinit's current strategy towards Linux, as well as its plans for the

future, *Linux Pro* spoke to Iain Cheyne, the Sales and Marketing Manager of Xinit to bring you this exclusive.

LINUX PRO: What would you say are Xinit's advantages over other server makers?

IAIN CHEYNE: First and foremost that would be our dedication to Linux as an excellent platform for running network services. We are 100% committed to Linux, unlike some of our competitors who'd be happy to sell any OS under the sun. Some of our customers do however run mixed Linux / FreeBSD networks, so we've now made it a point to qualify our systems with the newer versions of FreeBSD.

The other key factor that makes Xinit stand out from the crowd is the build-quality of our kit – it is nigh on impossible to find a difference in build-quality between an Xinit server and a branded server from the big-name server vendors. We achieve this by selecting only the best components from leading manufacturers. As a result, our systems tend to cost ever so slightly more than what you

would get from a 'whitebox' vendor, but then our clients appreciate that they are getting superb quality for much less than an alternative from big-name vendor. Then there is our automated testing process post-build; any components that are not up to scratch quickly stand out like a sore thumb and are replaced.

LXP: How eager is Xinit to take on board new technologies, such as Opteron and Itanium?

IC: Xinit Systems is always looking ahead to see how we can improve the IT infrastructure of our clients. 64-bit



The new SharqStor server is designed to provide high-speed data capacity through Xinit's custom interface - look for a review next issue.

computing is just starting to take a foothold in the market and Xinit Systems will shortly be releasing Opteron- and Itanium-based systems for both enterprise and scientific computing applications. Our approach is that clients don't need to go 64-bit unless there are compelling financial and performance advantages of doing so - we won't push 64 bits for the sake of novelty.

Then there is the issue of ISV support. ISV's are just starting to go gold with their 64-bit apps for Opteron and the third revision of Itanium - so at least customers have something to run on their shiny new 64-bit boxes.

LXP: To what extent is Xinit's server designs dictated by customer demand?

IC: Some of our competitors have more server "models" than there are stars in the sky. Xinit on the other hand, likes to keep things simple. There are only so many types of system configurations that are sensible from a pricing standpoint. We have a good range of systems that balance all the features - performance, availability and scalability - that one requires from a server. Usually, when a customer comes to us with a requirement that we do not normally offer, we first try and see whether there is a way to tailor an existing system to fit this requirement.

Sometimes this is not possible, so we work with the client to develop a solution that fits their requirement exactly. Naturally this process takes longer as we have to qualify the solution and ensure it is something that fits within our quality standards. What we won't do is bunch up a few components in a hurry just to make a quick buck.

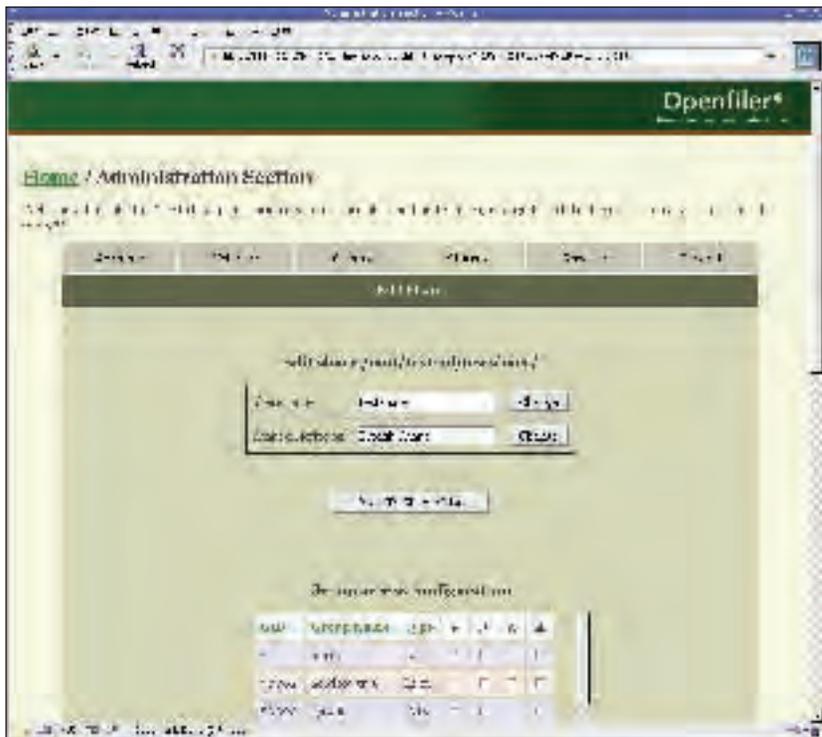
LXP: What is Xinit's strategy towards storage?

IC: While the company has traditionally focused on server infrastructure in terms of the hardware that we offer, there is a continuing shift towards developing the storage side of the business. So much so that Xinit Systems has developed a home-grown storage management interface, written from scratch, powered by the usual Open Source suspects and a few unusual ones. For instance, our storage management solution will provide iSCSI target, *ie* block storage, support from the same storage pool as file-level storage.

We have a storage management interface to deploy on our forthcoming SharqServ NAS appliances, and also intend to open source the management interface, called Openfiler, shortly after releasing the SharqStor range. SharqStor appliances are, for lack of a better phrase, a breath of fresh air.

In a market where customers are forced to choose from well-engineered but expensive products from the likes of EMC or NetApp, or settle for less expensive and consequently feature-challenged solutions from the smaller players, Xinit Systems SharqStor bridges the divide. The SharqStor range uses best-of-breed technology, featuring Serial ATA disks - reducing costs without compromising performance, scalability to 6TB in a single box, and, the real kicker: iSCSI target support.

Customers can now deploy a single appliance and have both block and file level storage exported from the same storage pool - killing two birds with one stone you might say. Our storage range also includes traditional SCSI JBOD and



The SharqStor management interface works entirely over the web, allowing you to create drive shares and configure authentication simply through your browser.

RAID solutions, as well as fibre channel-based solutions for clustering applications and high-end enterprise deployment.

LXP: Where do you see the server market going in the next few years?

IC: Simple answer - Blades. And this will be dictated by the Grid paradigm that is upon us. In five years, maybe even sooner, nobody will be talking about the Internet anymore - it will be called the Grid. Blades fit in well with the Grid paradigm because they are so easy to deploy and manage. You don't have to bunch up storage with the processors.

It goes something like this - keep CPU and memory together in the Blades and get the system people to look after it; get storage bunched up in NAS and SAN and get the storage experts to look after it; sell/export CPU and storage resources to partners/suppliers/customers either as a whole or individually.

Naturally, there will still be room for the standard rack servers we have today because there is some considerable work left to do on Blade standardisation. The interconnect will be Infiniband and all storage will be in SAN and NAS islands. There will be a resurgence of SSPs as they find a new life selling storage capacity to Grid-enabled enterprises. Oh, and it will all be powered by Linux! ■

"Xinit Systems is unlike other 'whitebox' server vendors in that we don't just throw a few chips into a piece of bent metal and call it a server."

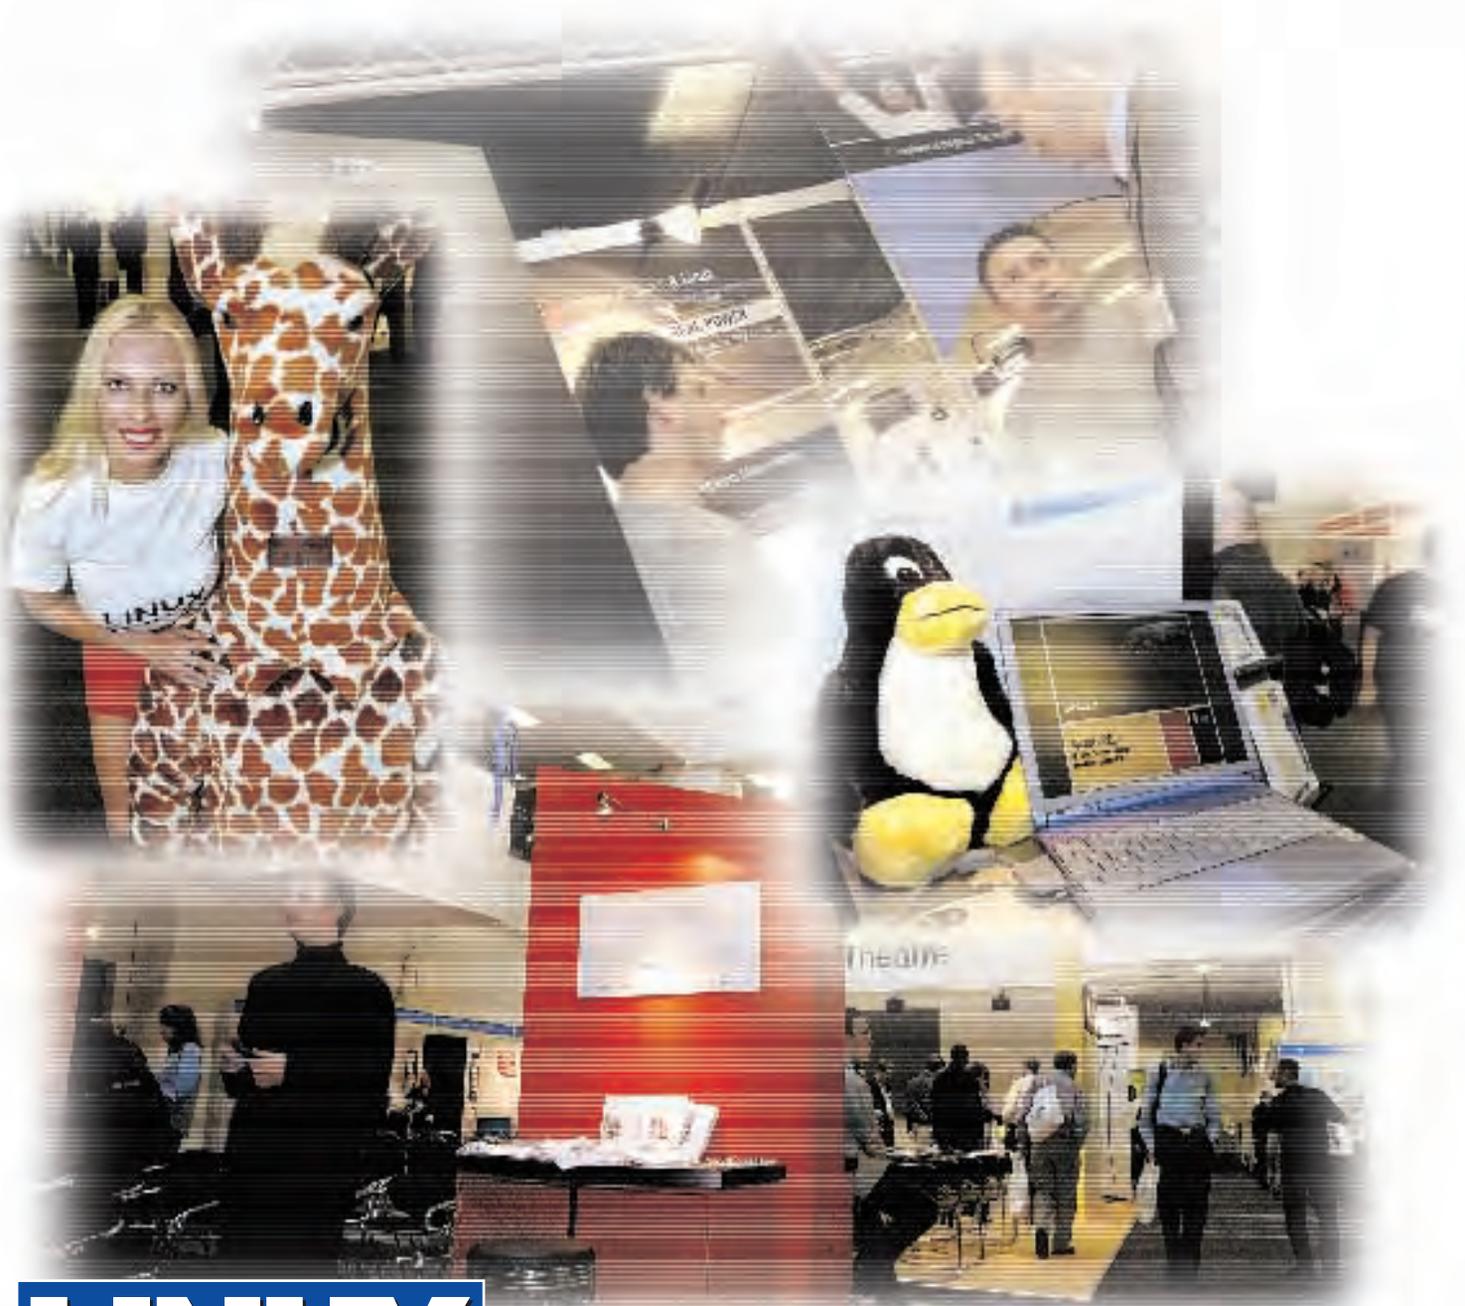
IAIN CHEYNE, XINIT SALES AND MARKETING MANAGER



LINUX EXPO 2003

OFFICIAL
SHOWGUIDE

8/9 OCTOBER, OLYMPIA, LONDON



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

PREMIER MEDIA PARTNER



Linux is in the on demand world. Can you see it?

Linux is everywhere. No other server operating system in history has spread as quickly across so many systems as Linux. IBM has more Linux-related offerings and the broadest line of servers that run Linux in the industry. Learn more about the IBM/Linux story and [@business on demand](#). Visit ibm.com/linux/seeit



IBM, the IBM logo and the e-business on demand logo are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. ©2003 IBM Corporation. All rights reserved.