

LINUX

THE UK'S BEST-SELLING LINUX MAGAZINE

FORMAT

TURN OVER
FOR YOUR
FREE
LINUX PRO
MAGAZINE

FASTER FILESYSTEMS

Maximum speed, ultimate reliability, infinite scalability – discover the right format for you



BOOK BONANZA!

Twelve of the latest releases plus our selection of must-have classics **p32**

DESKTOP DISTROS

Latest releases from Xandros, Slackware and Sun's Java desktop on test **p24**

BUILD BETTER CODE

SHOOTING FISH: moving sprites in our **SDL** game project **p78**

NEW **PHP** project starts inside – building a website from scratch

MASTER KDE/Qt signals and slots with **KDevelop** **p88**

SHARP ZAURUS

0980-75



Available on import – but should you buy? **p18**

ART ATTACK

Top vector art packages go head to head in 7-page roundup **p40**

MONO 1.0 OUT! **p06**

INTEL STRIKES BACK

New EM64T processors will join the 64-bit desktop fray **p62**

'Information synchronisation for the technical community' – POPsearch **p51**

DVD issue also available Printed in the UK
LXF57 SEPTEMBER 2004 £5.99
  0 9
0 771470 423010
£5.99 outside UK and 2nd class in Ireland

www.linuxformat.co.uk

Make a choice

Laziness, while certainly being the mother of invention, is not excusable when it comes to examining the options for making your life easier. If you always accept the default behaviour, life may be straightforward, but probably not very rewarding – but if you're reading a Linux mag I probably don't have to tell you that!

The point is that the 'default-acceptable-to-everyone' behaviour of anything, and in particular in the case of a Linux distribution, is the route of least resistance, common denominator computing which isn't going to get you the best results all the time. That's what we've explored in a small way this issue in our cover-feature. Practically everyone installing a Linux distribution these days chooses the default ext3 filesystem. ext3 is a great filesystem, there's no denying, but it isn't always the best one for the job. There are dozens of filesystems supported by the Linux kernel. These aren't just platform-specific options for compatibility like FAT and the Amiga FFS, these filesystems are usually developed with a particular

purpose, be that extra reliability, greater consistency, more robust multi-user use, or even speed. Even speed itself breaks down into different categories, depending on the predominant size and usage of the files.

Here, you can indulge your laziness, because we've done the hard part of the job for you: assessing the differences, creating and running the benchmarks. All you have to do is read the results. But as you do so, remember that as Linux is about choice, it's always a good idea to know a little about the choices you are making.

On that note, our *Roundup* this issue looks at the various software available for what is variously known as vector, scalable or structured artwork. A number of options are available in this fast-growing area of interest (because of SVG), and the best contenders are given a going over.

We never forget there is another choice you can make – whether to buy this magazine or not, so if you have any feedback on the magazine, please email or write to us (see details below).



Nick Veitch EDITOR



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.



Which Filesystem?
The choices – and why you would want to choose – examined and explained **p54**

Vector artwork –
another ten-gallon hatful of choices, but which is best? **p40**


What Open Source apps should you be installing? Hot Picks has the answer! **p48**



40



48



54

MEET LINUX FORMAT'S TEAM OF WRITERS...



Andrew Channelle
Newsman and newbies' best friend, Andy shows us that running web servers isn't as tricky a task as it sounds.



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



Mike Saunders
Like a big game hunter, he stalks the Internet's undergrowth to bag you all the best Open Source goodies.



Jono Bacon
Core KDE developer, web developer, sound engineer, freelance writer, musician – and ready for forty winks!



Paul Hudson
Currently packing his bags for OSCon, making sure to leave lots of space for Junior Mints, Jolt Cola and Mountain Dew on the return journey...

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

Hoyt Duff
Leading Red Hat Linux book author and fearsome scourge of mailing lists that get their facts wrong.

Biagio Lucini
Never one to miss a clock cycle, Biagio is intrigued by 64-bit computing in this issue's *What on Earth*.

Richard Drummond
Since moving abroad, he's had to cope with different filesystems just to survive...

Michael J Hammel
A professional artist, his regular graphics tutorial now covers vector artwork in *Inkscape* as well as *The GIMP*.

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 p13, p99 & p103

Contents



Welcome to another jam-packed issue of *Linux Format*, your guide to all things Linux! **LXF57 September 2004**

»» COVER HIGHLIGHTS



FASTER FILE SYSTEMS 54

Give your hard drive a new lease of life with our essential guide to filesystems

18 Sharp Zaurus SL-C860

Straight from Japan: the most advanced PDA around!

32 Books bonanza!


The essential collection for all serious Linux users

40 Roundup: illustration apps

Why your choice of art package matters...

76 Advanced Scribus

Get beyond the basics with our exclusive DTP tutorial!



32

'Look at Trolltech to see how a company can happily maintain a handle on development, even though an app is GPL...

08 We all want Sun to open-source Java!

'How long before the UK stops pouring our money down the proprietary drain, and lets children learn freely with FLOSS?'

06 PRO Parents & teachers - ACT NOW!

LINUXPRO

FREE EDUCATION!

Learn from the schools already using Open Source **1&1 AND LINUX**

Why its new web admin tool was developed on Linux

NOVELL & MONO The future of C# and GNOME

WINDRIVER Embedded

Linux has a new contender



18 Like a lovely little silver laptop – only smaller!



- 06 News**
Mono 1.0, SUSE, WineX 4, and more
- 12 Mailserver**
Share your views and experiences
- 18 Sharp Zaurus SL-C860**
Linux PDA with features galore
- 22 Northland**
Isometric control-freakism in motion
- 24 Slackware 10**
The power-users favourite
- 26 Xandros Desktop 2.0**
Business Desktop for the masses?
- 28 Visual SlickEdit 9**
A geek is only as good as his tools
- 30 Java Desktop 2**
X Windows within X Windows?
- 32 Book reviews**
Eight-page collection of the best
- 40 Roundup: illustration apps**
Five of the best vie for first place
- 48 Hot Picks**
The best Free/Open Source software
- 54 Filesystems**
Don't accept the default ever again...

- 62 What on Earth is...**
Intel's EM64T?
- 68 Beginners' tutorial**
Set up a web server in 20 minutes
- 72 Vector art tutorial**
Inkscape and The GIMP compared
- 76 Desktop Publishing**
Scribus – QuarkXPress-beater!
- 78 Games programming**
Squeeeeeeeeeeeal, fishies!
- 83 PHP**
Building a full site from scratch
- 88 KDevelop**
Just what *are* signals and slots?
- 88 Answers**
Your queries answered
- 99 Back issues**
Get the magazines you've missed
- 100 SUBSCRIBE AND SAVE**
Go on, you know it makes sense...
- 102 User Groups & advocacy**
Local support and fun
- 114 Next month**
Your guess is as good as ours...




06



40



62



72

Coverdiscs

A DVD or 3 CDs packed full of the latest Linux goodies **105**



» CDS A, B AND C

SLACKWARE 10 Linux's most venerable distro is still going from strength to strength
GCOMPRIS 6.1 Start them young – make Linux-based learning fun for ages 2 to 10
DIGIKAM 0.6.2 Manage your images
WEBMIN 1.150 Configure Linux systems from the comfort of your web browser
MOZILLA FIREFOX & THUNDERBIRD New stable versions of browser & mail client
STELLARIUM The night sky on your PC
ROSEGARDEN 0.9.8 Linux's premier pro general-purpose music composition tool



» DVD

SLACKWARE 10 Boots straight from DVD!
SUSE LINUX 9.1 PERSONAL Full distro ISO to install to your hard disk. Includes Internet, office, multimedia, games software, Firewall and Open Source YaST tool for easy installation, configuration and updating
KDEVELOP To build your own LXFGallery

Please read the coverdisc instructions on page 111 thoroughly before installing from coverdiscs!



SAVE MONEY!

SUBSCRIBE TO:
LINUX
FORMAT

See page 100 or phone 0870 4448645



Newsdesk

■ Linux to dominate supercomputing ■ Linux drama! ■ MIDI ■ Allied Irish Bank moves to Linux ■ SETI ■ Apple Rendezvous for Linux ■ Dell Linux desktop PC ■ Skype client ■ Acrobat Reader for Linux CE


Mono 1.0 is go!

The Open Source implementation of Microsoft's .NET framework released

Ximian has announced the release of the first official version of Mono, the Open Source implementation of Microsoft's .NET framework for use on Linux, Mac OS X, UNIX and Windows. Mono was said to be the primary reason Novell approached Ximian as a take over target last year, and the release is said to be vital to the company's future Linux ambitions.

The Mono project has been led from the beginning by GNOME and Ximian luminary Miguel de Icaza, who said the release of Mono should make it easier for programmers to create quality, cross-platform applications.

"Even as Linux grows on enterprise desktops, developing applications for the Linux desktop has been challenging because existing tools were extremely technical and complex," he said. "Mono is an extremely usable, commercial-



Dashboard attempts to provide information relevant to the task at hand, and was written entirely using Mono and C#.


grade development platform for Linux desktops and servers with a complete set of tools and APIs. It's based on

published standards and proven programming languages and libraries." This, he said, made it the perfect vehicle for developing applications for Microsoft Windows, Apple Mac OS X and other flavours of UNIX.

The package

The Mono package includes a C# compiler, a .NET-compatible runtime and a pair of application programming interface (API) stacks to deal with both native Linux servers and desktops and Microsoft's own .NET Framework.

A new website, www.mono-project.com/, has been launched to act as a hub for future Mono development, providing users with access to tools such as *MonoDevelop* (a native integrated development environment), support and the latest news on the project.



Muine is a music player that has been built using Mono.

Since the release of the first beta, an estimated 50,000 users have downloaded the software and, internally at Novell, it has been extensively used in the development of a number of projects, including *iFolder* and *ZENworks*.

With version 1.0 now in the wild, Ximian developers are aiming to keep pace with Microsoft's evolving .NET Framework. Mono 1.2 – which should see the light of day by Christmas 2004 – will include support for ASP.NET 2.0, System.XML and new compilers, while version 2.0, planned for mid-2005 will offer improved support for Windows.Forms and .NET 2.0 API.


The project was launched in 2001 and was seen by some as a validation of Microsoft's attempt to drive the IT sector into an ever-closer relationship with the MS way. Ximian took a lot of flak – and some derision – for attempting to work alongside the notoriously predatory Redmond giant. For more on Mono, see pages 18 and 19 of this month's *Linux Pro*.

SUSE 9.1 PERSONAL FOR FREE

Novell/SUSE has decided to once again offer a basic Linux distribution as a free download. In the past, the company has restricted free access to either Live CDs or FTP installations from SUSE mirrors, but this break with tradition offers downloaders the option of burning a complete version of SUSE Linux 9.1 Personal.

Personal is the stripped-down single-disc distribution but, after installation, users can change the installation source to point at SUSE's primary servers for version 9.1 and gain access to all but the proprietary elements of SUSE Professional 9.1.

Understandably, the product ships without any installation or ongoing support beyond what is available on SUSE's forums. LXF's DVD readers will find SUSE 9.1 Personal on this month's coverdiscs. For CD readers, the ISOs, and installation instructions, are available from www.suse.com/us/private/download/ftp/personal_iso_int.html and you'll need a blank 700MB CD to hand to contain the software.



ISOs available now. Get 'em while they're hot!

The war that never went away

Even MSN's *Slate* magazine recommends Mozilla!

The browser wars, written off as a one-sided battle won by Microsoft, are apparently not over yet. After another bruising month for the security of *Internet Explorer*, the US Department of Homeland Security (DoHS) took the step of recommending that IE users migrate to more secure offerings from the likes of *Mozilla*, *Firefox* and *Opera*. The Internet Storm Centre, which monitors worldwide Internet traffic to pinpoint sources of security problems also suggested alternatives to IE would make browsing safer. These warnings came as a result of the latest 'phishing' scam, which attempted to use a combination of vulnerabilities in MS's *Internet Information Services (IIS)* web

server and *Internet Explorer* to install trojan software, including a key logger, onto users' computers.

Step up *Mozilla Firefox*, a web browser now seen as so superior (even MS's own house magazine *Slate* recommends it) to IE that it has apparently inspired Bill Gates to reconvene the *Internet Explorer* development team in order to combat the threat. Insiders at [Mozilla.org](http://www.mozilla.org) say that following the announcement by the DoHS, which followed a similar warning from the US Computer Emergency Readiness Team (US-CERT), the Firefox server logs displayed a major spike in downloads. www.mozilla.org/products/firefox



Smart, secure, and blessed by the US's scariest administrative behemoth.

Like plugins, only different

Reducing Web monoculture by untying browser standards

The Mozilla Foundation has joined forces with Adobe, Macromedia, Apple, Opera and Sun to come up with a new standard plugin architecture designed to offer a rich web experience while adhering to the standards which make the Internet work. The extensions to the *Netscape Plug In Application Program Interface* will be implemented in the next generation of *Mozilla* and *Firefox* web browsers, as well as in future releases of *Opera* and *Safari* for Apple's OS X.

Mitchell Baker, President of the Mozilla Foundation said that plugins

were a critical component of the web experience, but that the technology needed to move with the times. The new initiative, she said, would make sure uncompromised access to web sites and applications wouldn't be restricted. "This will allow users to choose among a range of browsers without sacrificing interactivity."

What remains unclear is whether the plans will result in a cross-platform plugin architecture and the long-awaited arrival of native *Quicktime* and *Shockwave* players for Linux.

newsbytes

■ As part of its strategy to 'not be evil', Google has announced it will be 'opening up' some of its own code. Wayne Rosing, the company's VP of engineering, told a group of potential recruits that the time had come for Google to "give something back." It wasn't fair, he said, for a company to pick the cream of the crop (in terms of development talent) and keep the results of their work in-house. This would not, he said, be a case of open-sourcing Google. "That would be a little dumb when we have these Microsoft guys making noise," he said.



■ **Amarok** is yet another media player for KDE. The project has recently released version 1.0 of the software that combines iTunes-like access to local music, plus decent streaming abilities and the possibility of using GNOME's native *gstreamer* as its back end. Whether this or *Juk* will achieve 'default' status in the next major KDE release, LXF will certainly be reviewing it next month. <http://amarok.kde.org/>

■ **Slackware 10.0** has been released, and once again the maintainers of the distribution have shown their independent streak by shipping with a default 2.4 kernel. Slack 10 does include a 2.6 kernel for those adventurous users, but, it seems, the distribution is keen on retaining its reputation for stability over razzamatazz. The package does, however, include the latest releases from GNOME (2.6.1) and KDE (3.2.3) www.slackware.org/

■ The fantastic **Wikipedia** project has hit 300,000 English-language articles. The encyclopaedia's tally of over 90 million words easily outstrips its traditional rivals – including Encyclopaedia Britannica's paltry 85,000 articles and 55 million words. The biggest threat facing the project at the moment is a lack of developers with MySQL and PHP skills. The site itself is currently serving up an average of over 60 million pages per month and, as ever, is in need of cash injections both large and small. www.wikipedia.org/

■ For the first time in the history of the European Union, a sovereign state has withdrawn its vote of support for an EU directive. On July 1 2004, the Dutch parliament directed its Secretary of State for Economic Affairs to formally withdraw support for the widely reviled directive on software patents. With the Dutch authorities setting this precedent, it is hoped that other governments will take a critical look at the directive and follow the Dutch example.

David Cartwright

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



COMMENT

Ministry of silly forks

“The French Ministry of Equipment (MoE) is one of the latest Government bodies to replace its Windows installation with Linux – in this case Mandrake. www.mandrakesoft.com/company/press/pr?n=pr/business/2490. All is, however, not as it seems. Mandrake proclaims that: “With Mandrakesoft, the Ministry chose an alternative to proprietary software, and this is a significant step towards technological independence.” And were this true, it would be a good thing. Elsewhere in the press release, however, the company admits that “To meet the project's needs, two specific Linux distributions based on MandrakeLinux Corporate Server were developed.”

Let us get this straight: MoE is replacing Windows NT installations (presumably based on standard distributions of Windows NT that the entire world knows about) with a version of Linux which has been custom-built for it, and which therefore ties them into Mandrake to a certain extent. So by choosing “an alternative to proprietary software,” MoE seem to find itself using a specially built distro. Call me *Monsieur Picky*, but this sounds *très* proprietary to me.

There is a fine line between proprietary software and ‘standard’ software whose parameters have been customised and tuned for a specific customer. This project sounds like it sits the wrong side of that line. This whole story strikes me as a step backward for Linux – not just because MoE has chosen something that's not straight off the shelf, but more because Mandrakesoft thinks – and claims – that it's precisely the opposite.

Linux 'will dominate the Supercomputer sector'

The continuing lack of Microsoft's 64-bit mettle puts Linux's pedal to the metal

Latest figures suggest that Linux is currently running half of the world's 500 fastest supercomputers, but experts predict that our favourite OS is poised to dominate the sector absolutely in the future.


Analyst and Harvard Research Group Vice President Bill Claybrook maintained in an interview with Newsforge.com that Linux and Intel's fortunes appeared to be tied on the list. He said that Intel accounted for 119 of the machines a year ago, while this year's list includes 287 Xeon or Itanium systems. "That tells you right there that Linux is dominating," he said, as most of the machines were running Linux. Of this number, 243 are clusters, but Claybrook says the list will see another big shift next year, thanks to the price/performance ratio made possible by the combination of Open



Source software and commodity hardware. "Eventually, I think you'll find Linux is going to replace everything on the Top 500 list."

The fastest computer in the world – a record held since 2002 – was Japan's NEC Earth Simulator; while Linux was driving the number 2 beast, the 4096 Itanium2 processor cluster at the Lawrence Livermore National Laboratory.

In terms of vendors, IBM regained the crown it lost last year to Hewlett Packard by securing three of the top five places on the 'fastest' table.



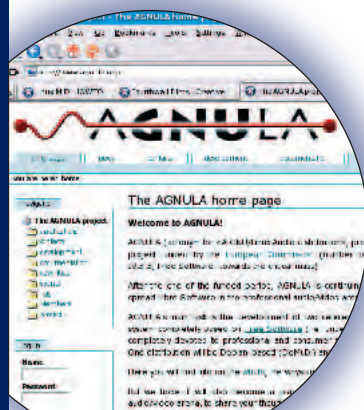
NEC's Earth Simulator retains its 'most powerful' crown for the third year.

LINUX WEB WATCH

Music to your ears...

Be entertained or get creative... or both!

MIDI, the Musical Instrument Digital Interface, has been exercising the LXF hive mind recently, in our quest to never have to reboot into an old OS again.



Our starting point on this trip was Phil Kerr's Linux-MIDI HowTo (www.midi-howto.com/), which runs through the fundamentals of getting ALSA installed, configuring devices and using Softsynths. There's even some useful stuff on latency there as well as a rundown – and links – to a wide range of MIDI players, sequencers and editors.

Once we started investigating beyond the information available from Kerr, we found the Linux Sound site (www.linux-sound.org/) which houses as near to an exhaustive selection of links to Linux musical applications as you'd ever want. We particularly enjoyed browsing the DJ

Don't go here expecting Abba's blonde chanteuse: that's *Agnetha*.

www.theora.org/ It's been slashdotted recently, so a Torrent or mirror site is probably your best bet.

software section: an under-explored sector of Linux audio!


Of course, if you want to get really dedicated to the glory of song, you'll want to check out AGNULA (www.agnula.org/), the best all-in-one distribution aimed at audio users. AGNULA is a complete installable distribution which is funded in part by the European Commission with the goal of bringing audio facilities to the masses. There are versions of the distribution based on Debian (DeMuDi) and Red Hat (Rehmudi). The distro includes solutions based on ALSA, JACK, JMax, Cecilia, the Csound effects system.

Finally, if you're more interested in listening and watching, and despair at



the current dearth of decent UK soaps, www.theora.org/ is offering streams of 'free content in a free format'. The highlight on the site so far is David Ball's 'edgy relationship drama' *Honey*, which is distributed under a Creative Commons licence. (<http://creativecommons.org/>)

newsbytes



■ **Element Computer** is taking a leaf out of Apple's book and attempting to tie its Linux distribution to hardware, which means everything should be optimised out-of-the-box, but with the cost benefits Linux can provide. Element's operating system solution is based on Xandros but also includes elements bought in from NeTraverse and a fully licensed copy of Windows 98, meaning applications such as *Microsoft Office*, *QuarkXPress*, *Money* and *Adobe Photoshop* will install successfully and run happily at native speeds. www.elementcomputer.com/

■ After a long wait, the powers that be in **Munich's** city hall have voted to officially adopt Linux as the desktop of choice for the city. Earlier reports had suggested that the project had fallen behind schedule due to both technical and cultural problems with the migration, but the council voted by a significant two-thirds majority to steam ahead with the plans. IBM and SUSE/Novell, which helped prepare the initial Linux bid, are favourite to get the final contract for the project, but this is not guaranteed.

■ **Red Hat** has GPLed its Global File System (GFS). GFS allows servers to read and write simultaneously to a single shared file system on a storage area network (SAN), thus improving performance and reducing complexity. GFS is said to be scalable to hundreds of servers and will work with any standard app. In addition to GFS, the entire infrastructure has also been put out under the GPL, including the CMAN cluster manager. Developers for the project needed! <http://source.redhat.com/cluster>.

■ The Search for Extra Terrestrial Intelligence (**SETI**) project has completed its transition to the BOINC (Berkeley Open Infrastructure for Network Computer) platform, releasing clients for the main computer architectures. While Windows users have a GUI for the app, users of Linux, UNIX, BSD, Solaris and OS X will have to make do with a CLI for the time being. Developers say the Classic SETI@Home client should be uninstalled before using the BOINC client as the latter has a lower CPU priority, but the original is still available.

■ **Apple** has made its Rendezvous technology, which offers zero configuration access to networks, available for Java, Linux, Solaris and Windows; and includes network protocols, identification and configuration of devices and services. <http://developer.apple.com/macosx/rendezvous/>

Allied Irish banking on desktop Linux

Securing the Punt for another Java Desktop System punter

In another corporate coup, Sun Microsystems has signed up Allied Irish Bank (AIB) as a client for its Linux-based Java Desktop System (JDS). The deal will see 7,500 desktops across the organisation migrate to Linux during 2005. The bank was lured by JDS's thin-client and desktop management tools and, a spokesman said, the financial advantages that Linux could offer.

Sun hopes that the deal with AIB will be a showcase for the banking industry, leading to wider adoption of Linux across the sector. Daniel Mayo,

of Datamonitor, told Computer Weekly that this could be a foot in the door for Sun. "A lot of banks are looking at refreshing their technology, which can be quite expensive," he said. "Linux is a good option."

Sun's Vice President of Desktop Solutions, Curtis Sasaki, said this was another example of the momentum that is increasingly building behind JDS and Linux across all areas of industry. "We are very pleased to have AIB as a client of Sun and see growing demand in the financial, education and government markets."

Dell gets Linspirational

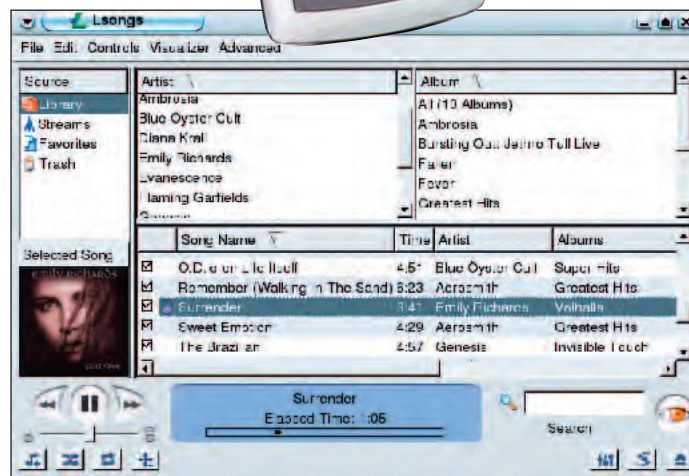

Strange 'support but not under warranty' situation

PC users in Europe can now buy a Dell computer complete with Linspire 4.5 and a year's membership to the Lindows/Linspire Click-'N'-Run warehouse.

Sold via Italy's Questar, the computers can be customised and shipped to any destination within the newly expanded European Union. The Dell/Linspire partnership in Europe was launched in tandem with the fully localised Italian version of Linspire and additional support in *ISongs*

(Linspire's own take on *iTunes*) for Dell's DJ MP3 player.

The range starts at 469 Euros (excluding local VAT) for a basic, monitor-less PC built around a 2.40Ghz Intel Celeron, 256MB RAM and 40GB hard disk. Delivery costs 30 Euros and hardware support will be provided via Dell's own support network, though the company was keen to stress that the software installation was not regarded as an 'official Dell install' and would not be covered by the company's standard warranty.



Dell's DJ jukebox is supported by Linspire's *ISongs* jukebox application.

Hoyt Duff

The co-author of *Red Hat Linux 9 Unleashed* runs a fishing pier when he's not being vociferous about users' rights on mailing lists.



COMMENT

What the future holds

“The development and growth of Linux is phenomenal. Just recently, it took significant technical expertise and a healthy dose of voodoo to enable dial-up networking. Nowadays, the process is essentially automatic – a non-issue.

The Linux community has grown by leaps and bounds, exhibiting a significant web presence: a Google search on the term returns 111,000,000 hits! Another positive indicator is the number and quality of publications focusing on Linux, such as *LXF*. I am astounded that the quality of documentation seems to be improving as well! Linux is more than ready for the desktop, the server... the world.

And what will become of beloved GNU/Linux? If we do it right, Linux will become a wholly ubiquitous, anonymous commodity. Its existence as a freely available, easily modifiable, and powerful operating system has it destined to become an ever present undergirdment of our world.

Is this a bad thing? Not at all. I remember when *VisiCalc* rocked the world: it was a marvel of an electronic spreadsheet. Alone, it was simply a novelty, but the creative use of spreadsheet software has helped transform the world. Emphasis on Linux itself will be refocused to what can be accomplished with it. The New Zealand stock exchange now uses a stripped-down version of Linux to run its Oracle database with fantastic results. Both Oracle and Novell see Linux as a commodity OS for their products.

The Revolution will have succeeded when we no longer need to think about Linux. Not long now...

Skype hyped for Linux

Cross-client capability still needed before mass adoption of VoIP can take off for High Street consumers



Skype brings VoIP to cross-platform chatters.

Voice over Internet Protocol (VoIP) is being hailed as the next big thing in the computer revolution. The ability to talk over the already established IP network that spans the world has the potential to make international calls a no-cost endeavour.

Skype, which uses the same peer-to-peer technology as the KaZaA file sharing network, is said to offer better-than-phone quality, full-duplex chat to users across the world. And now there's a beta client for Linux.

The package is available in RPM format for SUSE, Mandrake and RedHat/Fedora and Sun's Java Desktop System, and system requirements are a 400MHz x86 PC, 128MB RAM, a working sound system and, ideally, a broadband connection. The system allows users to talk to other Skype users without, currently, having to negotiate pages of ads, firewall problems, or constantly dropped calls; see the site www.skype.com/download_linux.html. We hope to bring you more reports and a review in forthcoming issues.

UPGRADE YOUR KERNEL!

A new security hole has been found which would allow a local attacker to bring down a system with just 24 lines of code. The vulnerability, christened **evil.c**, affects both 2.4 and 2.6 series kernels on x86 and x86_64 machines, but has been classified as 'low risk' by most vendors, as anyone wishing to exploit the hole would need shell access to a machine. Root access, though, is not required.

Kernel patches were released swiftly, and users/sysadmins are being urged to patch all affected machines. Andrew Morton, maintainer of the 2.6 kernel series, said the threat wasn't very serious. "Bugs wherein local users can lock the machine up are not uncommon, and local users have always been able to bring a machine to its knees anyway - say, by using up all the memory," he said. Morton also criticised the pair of programmers who discovered the flaw for not notifying the kernel team before releasing the exploit into the wild.

EMBEDDED LINUX NEWS


● Embedded Linux took another step into the mainstream with the release of the Consumer Electronic Linux Forum's first official specification and reference implementation for embedded consumer devices. This premiere release incorporates work from several of the groups working within the forum, expanding on a number of platforms supported by the original **CELF source tree**. Scott Smyers, Chair of the CELF Steering Committee, said the release of the specification was "a reaffirmation of the strong desire of major CE vendors to promote and use Linux as a major technology component of digital CE products."

● Spherex, one of the world's largest speaker manufacturer and supplier of the Xbox sound system, has signed a deal with Empower - a maker of **Linux PDAs** - to create a new range of audio products built on embedded Linux. The series will feature the same Optimised Surround Sound system used in the Xbox which will be coupled with a hardware player built on Linux technology from Empower. The devices will be licensed and released under a variety of brand names.

● Adobe has launched a new version of its **Acrobat Reader for Linux CE** software. This embedded version will first be seen on Sony's new Japan-only GPS system. The Sony HDD AV Navi device is intended to be a fully-fledged in car information point, so data, such as maps and travel guides can be uploaded to the device in PDF format for later access on the road. Adobe says it is in talks with many other manufacturers about bringing the technology to Linux-based mobile phones, set top boxes and personal video recorders.

● South Korea's Samsung Electronics Co has launched the first Linux-based **NAND Flash software solution for next-generation mobile phones**. Samsung's Linux NAND Flash memory software allows the NAND Flash memory to store code as well as data, and so should lower overall costs and reduce space requirements in mobile handhelds.

● For more on embedded Linux, see pages 16 and 17 of this month's *Linux Pro*, and see page 18 for a **Sharp Zaurus** review.



New to Linux, but established in the embedded world.

SCO NEWS

■ The wheels of justice move at a snail's pace, and this seems doubly true in the US. Of note recently, IBM's lawyers berated SCO's latest attempts at prolonging the case. Despite earlier claims in the press by Darl McBride and others that they had "tons, sackfuls" of evidence that IBM was illegally contributing code to Linux, the company had yet to show "one single line of code" to back up the claims.

■ In March and April this year, SCO management spent US\$2.5 million buying back 290,000 shares of common stock. Its directors have authorised the purchase of 1.5 million shares of stock on the open market over the next 24 months.

Catch up with all the latest from Pamela Jones at www.groklaw.net/

WineX 4 unleashed and renamed

Rebranding exercise ends confusion between Wine and WineX

TransGaming has released a new version of its *Wine*-based games product for Linux. *WineX* has been rebranded *Cedega* and, according to the company, now supports over 300 games out of the box. New additions

to the support lineup includes *Hitman: Contracts* from Eidos, EA's *Battlefield Vietnam* and *Star Wars Galaxies*. TransGaming also says the highly anticipated *World of Warcraft* massively multiplayer online RPG will

be supported from the moment it goes live.

Company co-President Vikas Gupta said *Cedega* marks a sea change in the maturity of the project, hence the new name. *"The new name, Cedega, is meaningful and reflects the significant growth that both the product and TransGaming have enjoyed over the last few years,"* he said.

Gavriel State added that he thought that the release represented a real *"milestone in game portability"* pointing out *Cedega's* support for the *DirectX 9.0* API as well as technologies such as pixel and vertex shaders. *"Cedega 4.0 also includes a new advanced inter-process communication architecture that can double the speed of games which make heavy use of Win32 kernel synchronisation primitives,"* he said.

Cedega is the first release from TransGaming to get a full European release. Customers and subscribers can now pay in Euros while the installer has been updated to include support for Dutch, Danish, Portuguese and German in addition to English.



Star Wars Galaxies comes to Linux via *Cedega* – the new name for *WineX*.

Opening up the Looking Glass

3D desktop paradigm to completely change our computing experience?

Project Looking Glass, Sun's next-generation 3D desktop environment for Linux, will have been released under an Open Source licence by the time you read this.

Looking Glass was created as a side project by senior engineer Hideya Kawahara, but is still at an early development stage. Sun has promised the project will be 'handed over' to the Open Source community *"once it's completed"*. The company has been under considerable pressure recently over its ambivalence on the question of Open Sourcing Java.

Kawahara, who has worked a couple of hours every day on the project for the last two years, says he was inspired by reports of the next-generation desktops from Microsoft and Apple, and set about creating something to 'advance Linux' at the same pace.

The roadmap for releasing *Looking Glass* as Open Source is not clear – but it's likely it will be GPLed, and

Kawahara says he is looking forward to seeing what other hackers within the community can do with the codebase.



Sun claims *Looking Glass* will change the way we interact with the desktop and not just be lots of fancy-looking eye candy.

Jono Bacon

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



COMMENT

Curiouser & curiouser...

“It came as a little bit of a surprise to many people when Sun decided to Open Source its futuristic 3D desktop environment *Project Looking Glass*; when it looked like a technology with some real potential, but a technology that would most likely be made available under a restrictive licence.

At a time when Sun appears to be getting pistol-whipped by just about everyone, it has managed to pull this 3D gem out of the bag to a joyous chorus of great PR.

In some ways, it could be said that Sun has actually made its life harder. *Project Looking Glass* could face problems due to its close dependency on what many consider will never be a free technology: Java. This could face particular problems if the software is to enter Debian; remember the problems with KDE's dependency on Qt in the days before Trolltech GPLed it – we could experience the same issues with *Project Looking Glass*.

It makes me wonder what exactly Sun is still trying to hold onto with Java. The organisation has come forward with a plan to embrace Linux, it has open-sourced Solaris, *OpenOffice.org* and *Project Looking Glass*, but Sun is still clutching Java with a tightening grip. If it were to open up Java and allow the community to continue developing it, Java usage would rocket in the Free Software world; I'm confident Sun could still maintain its authority on the code. We only have to look at Trolltech to see how a company can happily maintain a handle on development even though an app is GPLed. Come on, Sun – take the plunge and Open Source Java – we don't bite...”

Mailserver

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

★ Letter of the month

This month's winner receives a copy of *The Complete Linux Handbook 2*

Sticky situation

I have just been glaring at the sticker the maker had stuck on the front of my new box. One of the 'Designed for Microsoft' stickers that seem to come as standard with retail PCs. Of course I peeled it off but it left behind a mess – though not as much of a mess as the rubbish that had been put on the hard drive!

On the front of your brilliant magazine, you supplied a DVD that cleaned up the mess from the hard drive and replaced it with something decent. On the front of your Ultimate Linux Box in LXF53, you had a Tux sticker that would

cover up the other mess beautifully. Where can I get a sticker like that, or could you possibly supply one with a future issue of *Linux Format*?

Kim Giles, via email

Curiously enough, we were discussing the possibility of doing a *Linux Format* case badge only the other day. I'm not sure exactly what people would want though – Tux, presumably? Perhaps some of our more artistic readers might like to sketch out some designs for us?

Anyway, for reconfirming that as a good idea, you can have a copy of *The Complete Linux Handbook 2* for your trouble.



Simply snip around the outline and glue onto your PC to cover badges of other OSes...

Model-Driven Development tool

I know from my commercial experience as an Acceptance Test Analyst on large business systems that it is possible to specify a business system completely and unambiguously from a domain model consisting of a static UML class diagram with dynamic business use-cases defined in structured English in terms of the static model.

Such model-driven development from essentially a Requirements-level model will be the Next Big Thing in IT, as small teams of five to fifty Requirements Engineers will be able to design, develop and deploy major IT systems very quickly, with the concomitant cost savings and business advantage.

If the Open Source/GNU/Free Software community were to produce a robust tool that could understand

the model and automatically generate deployable client and server executables, truly this would be the killer app for Open Source in IT, putting Linux on every IT desktop and server; the business case for moving to Linux/Open Source would be absolutely compelling, especially for the public sector.

Needless to say, the Big Boys will evade producing such a tool like the plague, as it will render their proprietary ERP and database products obsolete, together with their associated revenue streams.

The downside will be that Requirements Engineering (with Acceptance Testing) will be the only game in town, with tens of thousands of IT programmers, systems analysts, test analysts, DBAs etc being made redundant. The question is, would this be politically acceptable to the Open Source community? As I firmly believe that domain-level MDD is

going to happen anyway, I would far rather see the Open Source/Free Software/Linux community lead the way than see any of the Big Boys cashing in on it.

So, how about it...?

Martin Lawrence, via email

Well, it's an interesting theory – but from here it doesn't look like the world will need less programmers. I appreciate that customised, business-driven applications account for a huge amount of software development, but there are plenty of applications that don't fit into this model.

Real Tutorials

I really like the magazine and find the tutorials section really interesting and useful. It's great to have so much accurate information on the kinds of things people can do with Linux. My only problem is that most of them don't really apply to me. I'm not a programmer, I don't run Linux as a server – I just like to run it as a desktop. I use plenty of the familiar software that probably everyone else does – *OpenOffice.org*, *KDE*, *Mozilla* and such other applications.

“We appreciate that some *LXF* readers aren't that interested in programming – or even servers – but we have to maintain a mix to please everybody.”

SUBMISSION ADVICE

WHAT WE WANT:

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

WHAT WE DON'T WANT:

- Technical questions – direct those to our *Answers* 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

What I really want is tutorials that show me how to get using this stuff better, or do things I hadn't thought of doing. There are plenty of specific magazines like this – like *Computer Arts* or the various movie-editing mags. None of these ever mention Linux, so there is plenty of reason for you to cover it. *The GIMP* tutorial is good and some of the beginner ones, but they are mostly too aimed at beginners. Jane Levers, via email

Thanks for your comments. We appreciate that some of *LXF*'s readers aren't that interested in programming, or even servers, but we have to maintain some sort of mix to please everybody. In terms of Desktop tutorials though, check out our mini *Scribus* tutorial this issue. If you have any suggestions for specific tutorials, let us know. Oh, and watch out for our cover feature next month...

British Broadcasting Confusion

There was an item about Linux on the television programme *Click Online* that was shown on the BBC's News 24 channel on Sunday 13 June 2004. In the report, it referred to the installation of a Linux distro on a PC with Windows already installed in the following way: *"to successfully install Linux on a machine that already has an operating system does require quite a high level of technical expertise that most end-users just don't have."*

On my system, I have Windows 98 and Mandrake 10.0 (recently



Although the BBC puts all its media streams online using *RealPlayer*, it has tried its best to support Linux clients – if you have problems, check out the help page at www.bbc.co.uk/radio/audiohelp_nix.shtml

upgraded from 9.1 with the CD-ROMs from May's *LXF53*) and I have never had any installation problems. I would certainly never describe myself as a person with *"quite a high level of technical expertise"*!

Since you wrote in your article in *LXF53* about installing Mandrake 10.0, there was no problems with a pre-installed Windows XP partition and that you are experienced journalists in the Linux field, could



unique customers - unique hosting

Low Cost Domains

.co.uk, .me.uk

£8.95 / 2 years

.com, .net, .org, .info

£8.95 / 1 year

Email Only Plans

A reliable, professional email service that's easy to use. Configure your email client or use web-based *SiteMail* wherever you are.

Email Starter

- 2 POP3 emails
- 20Mb Disk Space
- Unlimited aliases

£29.95 / year

Email Plus

- 5 POP3 emails
- 50Mb Disk Space
- Unlimited aliases

£49.95 / year

Shared Web Hosting

All plans are available on Windows or Linux platforms.

All our plans are built on 99.9% uptime reliability, managed security, the best tools and software out there, and Professional Support 24/7/365. Whatever your level, trust Hostway to look after your unique website.

Silver

- 100Mb Disk Space
- 1Gb Bandwidth
- 5 POP3 e-mails
- PHP/Perl/CGI
- Free Scripts

from **£7.95 pcm**

Gold

- 200Mb Disk Space
- 4Gb Bandwidth
- 10 POP3 e-mails
- PHP/Perl/CGI
- 1 MySQL Database

£14.95 pcm

Exclusive to Hostway...

Merchant Manager (NEW!)

Open for business! Hostway's fully customisable, merchant-ready online Shopping Cart system.

SiteControl - The latest evolution of our renowned Control Panel takes even more strain out of managing your hosting environment.

SiteBuilder - Create your own professional, customised website in minutes - no programming knowledge required.

All prices exclude VAT and are subject to change.

call us now on 0800 180 1880 • sales@hostway.co.uk • www.hostway.co.uk

domain registration • web hosting • resellers • e-commerce • dedicated servers

Australia • Canada • Germany • Korea • Netherlands • United Kingdom • United States

HOSTWAY
THE HOSTING COMPANY

READER TIPS

INSTALL

I was somewhat surprised when I read the description and install instructions in the *LXF55 Coverdisc* pages about *OpenOffice.org 1.1.1* to see up to six hours quoted for the installation – using a faster CPU than my own (Athlon XP 1800+). It took me less than five minutes to install from the source Tarball using a different technique!

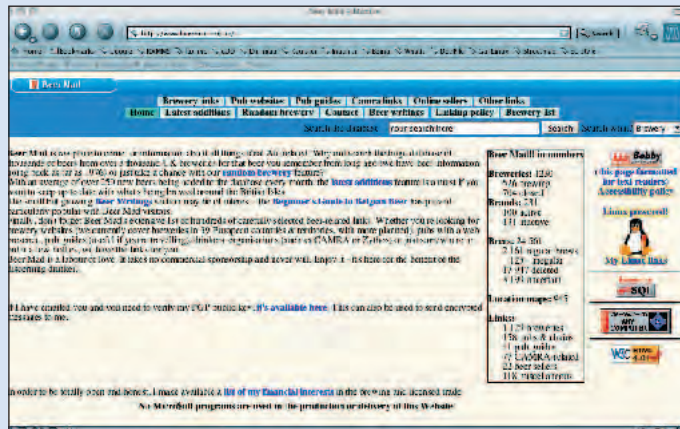
Obviously, I unzip and undo the Tarball, but add the *v* flag to the instruction and watch with interest as the 'SETUP' executable appears at the end of the files.

I then open the *OOo* install directory and head straight for the 'SETUP' executable, and follow the script which allows you to place *OOo* wherever you want it (*usr/lib*, *usr/local* or *home*) and within five minutes *OOo* is onboard. Amend the menu or set up links to *OOo*.

It has worked with all the previous *Star Office* and *OOo* flavours as it did with *OOo 1.1.1*. James Baldwin, *Sinnington, York*

LINUX & BEER

It's no secret that the binaries that come with your distribution have to be compiled in a very generalised way so they'll run on pretty much any x86-compatible processor. While recompiling code to take advantage of your particular CPU



Linux and beer, beer and Linux: indulgence in one can lead to inspiration in the other... See www.beermad.org.uk for a bit of both!

can result in performance improvements, it's true to say that most programs don't really use enough CPU time to make it worthwhile compiling your own. But, where you've got a program you use a lot that needs a lot of resources, the gains can be very worthwhile; especially if you supply compiler flags to optimise as much as possible for speed. I've been doing this for a long while now with *Audacity*, which I use a lot in my project of digitising my collection of over 1000 vinyl LPs – with superb results.

Of course there's no program any Linux user makes use of more than the kernel; for a while now I've been recompiling my kernel to

optimise for my Athlon XP 1400+'s instruction set, but I'd never got round to trying any further optimisations (or actually testing whether it was worthwhile) until a few weeks ago. To say the results have proved worthwhile would be like describing Vanessa Feltz as "not very shy."

I started off with a completely fresh 2.6.6 kernel bundle from kernel.org (as for some reason the latest Mandrake 10 kernel, 2.6.3-13, refuses to compile on my machine) and just made a few basic config tweaks to optimise it for my CPU. I built it and rebooted and sure enough, everything was working fine. Now I needed to time

something both CPU-intensive and reproducible – preferably without too much I/O as I didn't want any disk activity to affect the measurements. I settled for cleaning up a 16-minute music sample using *GWC* (Gnome Wave Cleaner) and timed the job at 1 minute and 7 seconds.

Now for the fun. I edited the kernel's Makefile and changed all the optimisation flags from **02** to **09** (the maximum *gcc* accepts, although I've since found out that everything above **03** is handled the same as **03** itself.) and rebuilt the kernel (as a new version, just in case...) Once it was rebuilt, I rebooted with some trepidation, because I've heard that kernels optimised above **02** are supposed to be unstable, but the machine came up OK and I could immediately see a big improvement in the time it took for KDE to load up. So how much faster would things actually run with the new kernel? I was expecting it to be significant but not huge; maybe 10-15 per cent faster, but when I reran my test, it only took *thirty-seven seconds!* In other words, almost a doubling of speed. Obviously not every program would speed up that much and it was hardly a comprehensive benchmark, but it was certainly indicative of the magnitude of the speed-up that I'd achieved.

« you please investigate and see whether the comments made are inaccurate. The website address is: www.bbcworld.com/content/temp/late_clickonline.asp?pageid=666&co_pageid=2

Alan, via email

It's nice to see the mass media paying some attention to Linux, but often coverage like this just reinforces stereotypes. Obviously, having to install Linux makes it 'harder' than if the computer came loaded with it in the first place! But it isn't really any harder to install Linux than it is to install Windows. The tack they followed in this piece was to install Linux on a pre-loaded Windows XP machine, which is probably as complicated as you are likely to get, but not really much of a problem if you use a distro that can

resize partitions (and for our money, Mandrake is still the best at this).

It does come down to opinion, but aside from being quite poorly constructed, this story doesn't really provide much real information.

Opera-ra-ra!

I love the mag, but what happened to *Opera* in the web browser roundup in *LXF55*? It's a great multi-platform web browser... was there a reason you left it out?

Godric Beresford-Jones, via email

Well noticed! Some other readers made this point on our forums. For reasons of space, we limited that particular *Roundup* to Open Source browsers, though in retrospect, it would probably have been fairer to include *Opera* too. Next time... [LXF](http://www.linuxformat.co.uk)



Just because a product has a 'community', this isn't an automatic indication that it is Open Source/Free software. If you've never wanted to see people with way too much time on their hands shamelessly exhibit their brand loyalty to a proprietary product, look away now...

But it gets better...

Discussing my findings next day at work (I'm in touch with a large number of Linux users via the company's private usenet groups) I realised that I'd missed out another acceleration trick I normally use; unrolling the loops (it's in the nature of C that it takes longer to execute a fixed number of iterations using a `for()` loop than it does to execute the same commands repeated the same number of times. So, when it's told to unroll the loops, GCC finds any `for()` loops with fixed iteration numbers and replaces them with the same number of repeated commands, which adds to the program size but makes it faster.) After recompiling my kernel with loop-unrolling, I reran my test and found that it had shaved another four seconds off the execution; not as big a change as the optimisation, but still a better than 10 per cent performance enhancement.

As I've seen so many comments about problems with kernels optimised above 02, I've been watching out for crashes ever since I

installed the new kernel. So far, it's been running for about a month without a sign of problems; my machine's running 24 hours a day with me doing all the usual desktop-type functions while I'm home, and the machine monitoring my website and running SETI@home and Folding@home the rest of the time.

My conclusion is that it's well worthwhile building a kernel with the maximum optimisation. Maybe I'd be a bit more cautious if it was in a production environment (the sort of place you'd use Slackware rather than Mandrake), but for a desktop machine there's everything to gain and nothing to lose. Just make sure you give your optimised kernel a name distinct from the existing one so if it does give you trouble, there's always the old one to fall back on.

Now I just have to find the time to try recompiling X!

Tony Green, Ipswich, Suffolk
www.beermad.org.uk/

RADEON DRIVERS

In response to 'Godawful Graphics' (Answers, p90 LXF55), I would like

to point out that ATI *does* provide a driver (it's available from www.atitech.ca/support/driver.html) for Radeon 8500 series and above under Linux.

I have found this to be an excellent driver for my Radeon 9200SE, and have no complaints about its 3D performance in UT2004. However, I did find that it wouldn't compile against the Fedora Core 1 default kernel, however using a 2.4 series kernel compiled from sources at kernel.org worked fine.

Richard Morgan, via email

LIVE DISTRO

I received LXF54 recently and noticed that it was featuring a review of live distros. I just wanted to say that I was disappointed that Mepis (www.mepis.org) was not included in your Roundup, as it is a great live distro that is extremely easy-to-both install and use. As a Linux newbie (Linux user/learner since November '03; LXF reader since LXF46 and a subscriber since LXF51), I have tried half-a-dozen distros and found Mepis to be one

of my favourites (I'm trying Fedora Core 1 from LXF49 now). I just wanted to see it receive some attention, as it's a great distro!

Scott Johnson, via email

There are many, many live distros and it would have been impossible to cover them all. Thanks for your tip though, we'll check out this one for next time.


KPACKAGE

I read the letter in LXF55 about Kpackage with interest, and I'd just like to point out, in fairness to the Kpackage developers, that it does support Debian packages, will do `apt-get update, upgrade, fixup` etc; and is a very readable, easy-to-use way of viewing the details and file list of a particular package. You can specify `apt` sources as well as local folders, and although I have noticed a couple of apparent display bugs, they aren't that serious – with a little more work it could be a worthy rival to `dselect` or `synaptic`, especially as it's so well integrated into KDE (I can't wait for the Woody backport of 3.2.3 in July 2004!)

Andy Websdale, via email

Helpdex

shane_collinge@yahoo.com



Reviews

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

LXF VERDICT EXPLAINED

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

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

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

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

Value for money/Documentation: Whichever is most appropriate!

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

●●●●●●●●●●
10 The close-to-perfect product

●●●●●●●●●○
8-9 Good, but has a few niggles

●●●●●●●●○○
6-7 Does the job, but needs work

●●●●●●○○○○
4-5 Average.

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

THE TOP STUFF AWARD

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



WHAT'S NEW...

Sharp Zaurus SL-C860 PDA

Faster than a speeding metaphor and more silver than a magpie's nest, but does the high price-tag outweigh the coolness factor of having Linux in your pocket? **p18**

Northland

Men are from Mars and women are from Venus – where they apparently spend quite a lot of time playing micro-management sims... **p22**

Slackware 10

Long-lived and battling on: will Slackware 10 bring back the glory days for this venerable Linux distro? **p24**

Xandros Desktop: Business Edition

Version 2 now comes with *StarOffice 7* and Codeweavers' *CrossOver Office* **p26**



18

Visual SlickEdit 9

KDevelop missing some of your most-wanted gadgets? VSE8 might just be for you **p28**

Java Desktop

Now 50% more double good **p30**

Books bonanza!

Contrary to *Postal 2*'s "save a tree, burn a book" philosophy, we review a bumper crop of learning literature and also pick our essential bookshelves for readers of all levels **p32**

LXFBENCH 2004 EXPLAINED

To comprehensively test the capabilities of machines we review, we have developed *LXFBench 2004*: a new benchmark suite designed to push hardware of all shapes and sizes to its limit.

The test is broken down into four distinct parts: multiprocessing, uniprocessing, RAM, and hard disk, of which the first two are largely similar. The multiprocessing test creates four child processes in order to take advantage of SMP hardware, then proceeds to run *oggenc* to encode a large audio file to Ogg format, uses the *GD* image library to resample a complex image several times, and also runs an external C program to calculate

the hashes of random numbers. The *uniprocessing* test is identical except that it runs on just one CPU.


Both the *RAM* and *HD* tests use the *SQLite* database library to manipulate database information in RAM and on the hard disk respectively.

The overall score is an average of all four tests, and is presented as a bar graph for ease of reading. A score of 1 means that the machine has equalled our yardstick machine – a 1.8GHz Pentium 4 with 512MB RAM and an IDE hard disk. A score of 2, therefore, means that a machine has completed our tests twice as fast as the benchmark. The majority of the code was written using PHP 5, with the CPU-intensive tests written in C.

BENCHMARKS

CPU	6.07
SINGLE	3
RAM	2.17
HD	0.46

OVERALL 2.93



All our benchmarks, unless specifically noted otherwise, are run on a fresh installation of Red Hat Enterprise Linux 3 AS for the specific platform. All source code, including PHP itself, is compiled using GCC unless otherwise noted. The *mhash* library, created by Nikos Mavroyanopoulos and Sascha Schumann, is used for data hashing.

LINUX-BASED PDA

Sharp Zaurus SL-C860

Along with sashimi, *Battle Royale*, and Sony's Aibo robo-pooch, the Zaurus is one of Japan's hottest exports – but it's a mixed bag, as Paul Hudson found out...

BUYER INFO


Grey-import PDA from Japan with Linux at its core. Also consider Sony's PEGUX50 or Dell's Axim line.

- **SUPPLIER** Shirt Pocket Ltd
- **PRICE** £565 (inc VAT)
- **WEB** www.shirtpocket.co.uk/

Many people have a hard enough time getting Linux to work perfectly on their laptops, and most will be content with a few things (such as IrDA or a built-in Bluetooth card) being disabled rather than spend the rest of their lives fiddling with options. How much harder, then, is it to get Linux working on a PDA? This is the unique selling point of Sharp's Zaurus PDA: get all the power of a Linux box crammed down into a unit that fits neatly in the palm of your hand.

Sharp and the UK

Sharp dabbled briefly in the UK PDA market with the release of its 5000 and 5500 Zaurus models, both of which were traditional, portrait-aspect PDAs that received a lukewarm




The folding clam-shell design means you get a great screen without compromising on keyboard space – this is the best PDA design out there.

reception from users in Europe. Since then, Sharp sold up its inventory of those models and nothing new has been released.

Back in its home town of Osaka, Japan, Sharp's Linux-appreciating research teams haven't been quiet, and a wide selection of models have been produced and sold over there

that build upon lessons learnt from the early models. This latest release is the culmination of Sharp's efforts, and is – sadly – not likely to be released over here. However, as we're rapidly becoming an international magazine, and also because an enterprising group in the UK has taken on the job of shipping these models in, translating



Qtopia is a pretty good attempt at a windowing system, but it could be improved so easily...

them, and selling them on in the UK, we couldn't wait to get our hands on it.


Before we continue, it's important that you fully understand that the model reviewed here is a Japanese product, with a mix of Japanese kanji and katakana on the keyboard as well as normal Latin characters, and a fully Japanese manual. Shirt Pocket is selling the translated version in the UK, but Sharp does not (and probably will never) officially support it here.

Unit design

The portrait layout has now been dropped in favour of a clamshell design similar to the Gameboy Advance SP, which means there's more real estate available for the screen and the keyboard. The touch-sensitive screen remains, but it's now a full 640x480 resolution – it might not sound much, but given that the viewable area of the screen is just 8cm

ZAURUS AT WORK

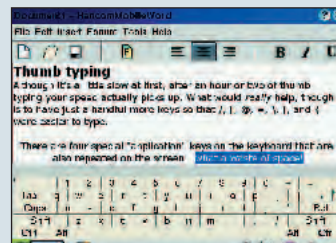
How the bundled applications help redefine the meaning of "portable office"




Manage your time and meetings with the Calendar application – it's fast, powerful enough, and makes good use of the small on-screen space.



The World Time program lets you point anywhere on the map with the stylus to get the time there, and takes into account seasonal variations like BST and GMT.



Although the keyboard could be better designed, your thumbs adapt soon enough and quite frankly anything's better than the handwriting recognition...



Hancom Mobile Sheet is all the spreadsheet most of us need, and even saves in Excel format, which is easily viewable by us Linux-using types in OpenOffice.org.

ZAURUS AT PLAY


Games, music, movies, and more...



Chess: the game where champions pit their skills against each other in an unforgiving battle for supremacy. Now, who goes first, again?



The Movie Player is capable of playing back most common formats, but bandwidth ends up being the bottleneck: we tried playing back an MPEG-4 movie from a Microdrive, and it dropped about every eighth frame.



Rocket Elite is a clever use of the touchscreen – simply prod the screen where you want the ship to fly, trying to compensate for inertia along the way... BOOOM!



The Doom demo is somewhat playable. If you have ever played *Doom* or anything else requiring a modicum of hand-to-eye co-ordination on the Gameboy Advance and found it hard to control, it's best to skip this one!

wide by 5.5cm high (10cm diagonal diameter) it's actually a very high resolution and graphics look fantastic onscreen. There are six levels of brightness for the backlight that go from quite dark to eye-burning bright, which means you've got maximum control over how long the battery life will last, but it's a shame there's no way of turning the light off completely.

The unit itself is 12cm wide by 8.5cm deep and 2cm high, which is a little bit bigger overall than previous clamshell designs released in Japan. The reason for this is that Sharp now bundles a longer-life battery that

actually sticks out from the bottom of the unit, warranting a false bottom to house the battery. It fits nicely into the rest of the unit, though, and the extra battery life makes a big difference.

Full keyboard

The keyboard is a full QWERTY, with numbers along the top row, cursor keys at the bottom-right, and some symbols scattered around. Most keys have extra functions that can be accessed through Shift or the Function key, and some even require you to have Shift *and* Function held down to get the key you want. This makes typing anything longer than a short note to yourself – without a hard surface on which to rest the unit – a bit of an abortive task. Ideally, Sharp would make some of the keys larger (particularly the left-hand Shift, Enter, Function, and Spacebar keys) so that moderately fast typing could be achieved. If you want a particular character without having to hunt through various combinations, there's an on-screen keyboard where you can just type with the stylus.

Failing that, you can try your hand at the handwriting recognition, although we wouldn't recommend it.


On Palm PDAs, you can Graffiti; which, although it takes a little time to learn, ends up being a very effective way to input data. On Windows PDAs, you can literally write on the screen and it picks up what you write with remarkable accuracy. On this PDA, however, you can write in perfect copperplate print and it will invariably interpret it as "7%\$@^^". We managed to get it to recognise most letters in the Western alphabet, but only if we inputted them very slowly and carefully – hardly an efficient way to work.

For connectivity, there's a single CompactFlash slot and a single Secure Digital, which might sound good until you realise there are no onboard peripherals. Devices that allow WiFi, Bluetooth, Ethernet, modem connectivity, screen output to a monitor, and even a digital camera all come separately and plug into the CF slot. As there's only one such slot, you can only have one in at a time. Having the SD card there does alleviate the problem a little, as at least you can use that for your storage and keep the CF card for WiFi, but Hitachi's Microdrive only comes in CF format, and at 4GB is the largest and cheapest storage format around. Having to choose between a Microdrive and WiFi is pretty hard, but what we'd rather see is onboard WiFi and Bluetooth so that the CF slot is effectively free for more storage.

The devices cost extra, and most of them aren't really worthwhile. The WiFi access is admittedly very easy, and we got a fairly good signal in our tests. As the screen is quite a high resolution for the small space it's amazing how well the text on websites comes out, and thanks to the landscape view you don't need to scroll around as much as you would on other PDAs. The screen output does work, but at a very slow pace: you type, and it responds a second or two later. Furthermore, 640x480 just doesn't look good on any monitor larger than 14". The digital camera is a waste of money, and should be avoided – the camera on your phone is almost certainly much better. There is a built-in speaker that manages to kick out quite a bit of noise for such a small thing, but there is no built-in microphone – only a slot for it.

Battery life depends entirely on how you're using the device, and will go from under three hours (playing the *Doom* demo at high screen brightness) to about a week (using it for under an hour a day on low screen brightness).

CHOOSE A VIEW



With the lid closed, the screen is safe from dust and scratching.



Clam-shell mode is best for typing and general use.



But rotate the screen around by 180 degrees and...



...the Zaurus looks just like a normal, portrait PDA.

Charging isn't particularly fast thanks to the low-power charger: we left it on for about five hours to charge, which seemed to do the trick.

It runs Linux...


Although there are some particularly smart PDAs out there, this is one of the few that runs Linux, which is what




Several keys retain their Japanese symbols, but it's easily ignored.

Zaurus AND LINUX


Just how hard is it to get to the core of the OS...?




The font is a little teeny, but you can just about make out the tabs at the bottom of the screen that show we've got several terminals running. It's *Bash*, so who cares?!



The only good GUI filebrowser on the system is *Tree!Explorer Qt* – the others mask the underlying Linux filesystem, limiting the scope for enhancing your geek cred!



The System Info screen offers the most raw information about the Linux OS and some associated applications for those who are really interested.



The *WhatsUpOnZ* system manager for processes, modules, CPU and Memory is like *top* (which is unavailable), and lets you kill rogue apps with the press of a stylus.

really makes it special. To save power, it's a highly specialised version that doesn't use X or KDE. Instead, it's driven by Trolltech's *Qtopia* window manager that uses *Qt* for its toolkit and works quite well. This is quite an advantage, because Trolltech makes the *Qtopia SDK* freely available to developers under the GPL so they can write applications for it.

...but it's well-hidden

Although the actual Linux side of the system is pretty well hidden, there is a "Konsole" link under the Applications tab that takes you straight to *Bash*. The default suite of terminal programs is pretty minimal, but you do get FTP, *Telnet*, *find*, *grep*, *more*, *ifconfig*, and even *vi*. There are some things missing, though, most notable of which is *SSH*. We consider *SSH* a critical part of any Linux system today, and it's a real shame it's missing here. Although you can search the web and find a download of *SSH* online, it's not as easy it as it should be – we'd like to see tools like *SSH* supplied on an CD ROM and easily installed by users who want them. I think most people would rather sacrifice the installation of *Netfront* (you get two web browsers on the device) and put Perl, PHP, or Python in its place!

The selection of GUI applications provided is good enough – you get *HancomOffice*, *Opera*, a few basic apps (text editing, etc), media

players, and a selection of demo games from Eon Games. *HancomOffice* does a pretty respectable job of importing *Microsoft Office* documents, and both the spreadsheet and word processor apps are powerful enough for the vast majority of users – not that we would recommend typing long essays with your thumbs!

Qtopia allows you to multi-task easily, and we could happily run seven or eight applications without any apparent slowdown. The interface does need a lot of work, though – it's quite ugly for the large part, doesn't give you much control over the colours of things, doesn't look good with any wallpapers, and everyone we asked to try out the program launch menu found they ended up selecting the wrong thing at least half the time. Come on, Trolltech – we know you can do much better than this!

Lasting impressions

While there's no doubt this product is very good, it's held back by several fundamental problems:

- 1 The CPU isn't up to the job. Intel now has a 624MHz X-Scale chip out with MMX that reportedly increases overall system speed drastically. We expect it to be about 70 per cent faster across the board than the 400MHz chip here, rising up to about 100 per cent faster for multimedia apps.
- 2 There needs to be more Linux on the machine. There's a huge market in the world for geeks who want *SSH*, *GCC*, and *Emacs* on their PDAs, and a really cut-down release of these extra features should take up no more than 10MB of system memory. At the very least, these could be provided on a CD so that they can be installed if users want to sacrifice the space – a technique used by Sony Ericsson with its mobile phones.
- 3 It needs WiFi built-in. Giving up the sole CF card slot just to connect to a network is unacceptable in today's connected world.
- 4 The price is shockingly high.

Someone, somewhere is taking a huge cut, because

this thing is simply not worth £565. You can get a 624MHz Dell Axim PDA with built-in WiFi and Bluetooth for just less than half the price of this unit, and that includes paying the Microsoft tax for the Pocket versions of Windows and *MS Office*. What's more, the manual and all the keyboard are in English!

5 The keyboard isn't great, but at least it's better than the dreadful handwriting recognition.

On the plus side, the screen is the best we've seen so far (although it would be nice if it used the full space on the top panel as opposed to having such a large bevel), the battery life is good enough, having a *Bash* console to hand – even if it's just with *Telnet* rather than *SSH* – is great, and the rotating screen means you get the best of both worlds.


Although it's a bit much to hope that all five of the problems will get magically fixed in the next release, we'd probably happily pay the £565 if Sharp could get the new CPU in, add *SSH*, and have WiFi in as standard. Until then, this remains little more than a pocket toy – albeit an expensive one – that has too many drawbacks to make it worth considering unless you're flat-out desperate to have Linux in your pocket. **LXF**

LINUX FORMAT VERDICT

FEATURES	7/10
PERFORMANCE	7/10
DOCUMENTATION	5/10
VALUE FOR MONEY	4/10

At half the price it would be worth it, but with a slow CPU, lack of built-in WiFi, and a high price tag, this is one to miss unless you just can't live without the screen.

RATING 6/10



Extra devices galore, but they all come at a price: talk to your bank manager first!

STRATEGY GAME

Northland

Point, click, point click, point click, poi-sdfsdllfkjsdffffffffffff. Paul Hudson falls asleep at his keyboard and gets his wife to do the bulk of the work...

BUYER INFO

Isometric micro-management sim. We would recommend also considering *Majesty*, (LXF48) or *Knights & Merchants*, (LXF50) but they're almost the same game...

■ **DEVELOPER** Linux Games

Publishing

■ **PRICE** £30

■ **WEB**

www.linuxgamepublishing.com

Some years ago, the very first *Settlers* game was launched and did rather well. The formula was quite simple: you trained soldiers, made weapons for them, then fought your enemies and won. Unfortunately, that series descended into ever more complex micro-management – weapons needed to be put together using wood and metal, food needed to be put together with wheat and water, and so on, until in *Settlers 4* you're essentially playing with a massively complicated maths engine. If you read our review of *Knights & Merchants*, you may remember that it was for the large part simply a clone of the *Settlers* series and its formulae, concentrating on the micro-management aspect of the game



Welcome to Smallsville, population 20. The villagers have rejected my plan to build a Dwelling 3 and instead have asked for a multiplex cinema and a ride-through fast food shack. They'd get what's coming to them if only building a guillotine didn't involve three hours' work.

rather than actual original storyline and new interaction.

Hot on the heels of *Postal 2* (8/10, LXF56), LGP brings us Northland – the sequel to the vaguely successful *Cultures 2: The Gates of Asgard*, and, rather sadly, yet another *Settlers* clone. Even more sadly, the game makes little attempt to alleviate the dreary detail of the genre – if anything, it increases it.

Pointy, clicky

OK, so it's your job to look after a Bjarni, who, despite his name, isn't a large purple dinosaur. Instead, he's the hero of a small Viking tribe, and you need to help the tribe grow whilst also aiding your friend, Hatchie. To do this you need to build houses where your people can settle, build places of work for them to spend time in, build barracks and train soldiers, and eventually fight your way to victory. So far so good, right?

Now, there's a problem: how do you get soldiers? First, you need to build their weapons. This requires a carpenter's shop and a carpenter, which in turn requires woods and a woodcutter. Second, you need the actual soldiers, which are trained in barracks. And how do you get barracks? By building them with builders. Builders themselves are just specialised civilians, and you need to get those civilians by finding two other civilians (male and female), commanding them to marry, building them a house in which to live, then ordering them to reproduce (you even have to choose the sex of the baby). After a time, a stork appears with a baby, which then spends a few


minutes crawling around before turning into a toddler, then continues growing up until eventually you have a person that can in turn reproduce.

Yes, you read that right: you need to personally instruct each member of your tribe to reproduce, otherwise they just kill time however they see fit! Similarly, you need to instruct them to do a variety of other basic tasks: otherwise they will happily sit around and do nothing – builders would rather sit and chat to other builders than walk over to a wood pile and pick up the necessary wood to do their job; a curiously accurate parallel of real life. Instead, you need to assign carriers to move things to and fro in your production line-like village.

The buildings themselves are no source of excitement. For example, you start off training a civilian to be a master farmer. Then, once they have achieved that, you build them a mill to grind the wheat, and it takes some time for them to become skilled as a miller. Then you build a well and a bakery so that, at last, your people can eat, and your miller becomes a baker. After a while being a baker the fellow becomes skilled enough to become a brewer, so you build a brewery and move him there, and so on, and so on. To make life more interesting, you can also upgrade buildings – for example, you can house your villagers in a 'Dwelling 1', but you can also choose from a 'Dwelling 2', 'Dwelling 3', or a 'Dwelling 4', each of which is successively larger. Only through hours of hard work and chasing around people to get them to reproduce can you complete a level and move on to the next thrilling plateau.

Click, click

While you might imagine there is only a limited number of ways in which a game could contrive to make you do the same thing again and again, you'd be wrong – after almost a solid day of playing this game (you'll never know how much willpower it took to do that) we found we still had to tell people to wear shoes (to make them walk faster!)



Micro-manager's log, star date 3106.04: In a desperate attempt to interest myself, I have chosen Bjorn for ritual sacrifice to some wolves. His wife, Ragna, and his five children will accept the fate I have decreed for him, otherwise they too will meet their doom. Mwahahahaha! Mwahaha! Mwaha. Ha. Oh, the hell with it. I wonder what's for dinner...

Xandros Desktop 2.0 Business Edition

BUYER INFO

Desktop optimised for integration in established Windows-based organisations. Competition from SUSE, Sun, Mandrake and now Red Hat.

■ **DEVELOPER** Xandros

■ **PRICE** US\$129 single user.
US\$495 Five licence pack


■ **WEB** www.xandros.com

Enter Xandros, and its award-winning Business Desktop product: sporting not just a selection of



The file manager deserves a mention for the extensions it adds to *Konqueror*, especially in the field of network access. As a client, this is by far the most elegant *Samba* solution we've encountered. The left pane of the file browser contains a standard-looking 'tree view' of the file structure, except the entries are broken down





Should a business need them, a range of Windows applications are supported at speeds almost indistinguishable from their native OS.

into logical parts such as 'My Computer', which equates to /home, 'All File Systems', which equates to / and is available only to the system administrator, and Windows Network. Opening this tree displays a link to every available SMB-based network; opening that reveals every shared resource, including directories and printers. Once a directory is selected, mounting it locally is merely a case of hitting the 'mount' link and providing a path to a location. The default is ~/user/[original name of directory]. The share is immediately available and automatically connected on subsequent boot ups. Booting up away from the network sees a dialogue pop up offering to permanently remove the share entry or leave it for next time. This same system is also usable for NFS shares, meaning a Xandros machine should drop into any heterogeneous network environment without difficulty.

All-round office suite

As is becoming customary, Xandros goes for overkill in the office suite stakes, adhering to the maxim that "choice is good", especially when dealing with 'better the devil you know' businesses. A key selling point appears to be the aforementioned *StarOffice 7*, which has both advantages and disadvantages over its Open Source stablemate. On the plus side, the spell-checker and thesaurus are much better; support – if needed – is available from a corporate-friendly source; and the individual apps launch fairly quickly. On the other hand, *OpenOffice.org* looks better, especially

in its Xandros guise: it is more up-to-date in terms of bug-squashing, and the font rendering in *StarOffice* is still in dire need of anti-aliasing.

For those who really fear change, or who don't want the training hassle of swapping office suites at the same time as operating systems, *CrossOver Office 2.1* (COO) handles legacy applications beautifully. As you'd expect from this superb application, *Office XP*, *Adobe Photoshop*, *Dreamweaver* and *Flash* installed and ran well and the integration of the technology with the OS is very pleasing: for instance, accidentally attempting to launch *Flash MX* that had been installed on the Windows partition of our test machine didn't cause an error; it just asked for the serial number and then launched. We did find, however, on one attempt that the install routine for *Office XP*

failed to add application icons and entries to the Start menu, and that occasionally the *KPanel* would disappear on running *FlashMX*, which simply required a desktop refresh to solve. This latter problem, though, seems to be a bug in COO rather than Xandros' implementation.

The developer's choice of applications is, on the whole, good and the ability to quickly add official (or even testing or experimental) Debian sources to the Xandros Networks (XN) package management system ensures that keeping the OS and applications up-to-date is a breeze. The decision to go with *Mozilla* for web browsing and mail is sensible in terms of visual integration and for the best performance with *CrossOver Plugin*, which is also included in the package, but it's nice to see both *Evolution* and *KMail* available in XN. This might become especially important if Xandros can quickly add *Connector for Evolution* to the system. It would then be a viable option for businesses tied, for whatever reason, to Windows/Exchange Server in the back room.

Conclusion

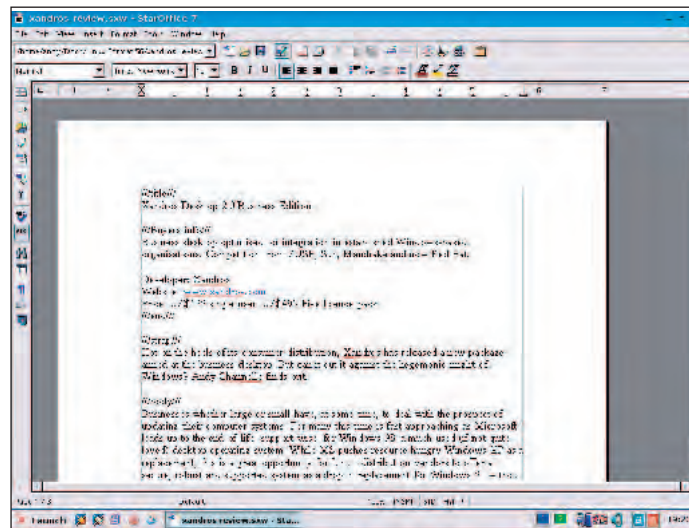
The prospects for desktop distributions appear to be strengthening, especially with Windows 9x coming to the end of its life, and Longhorn a distant but growing speck on the horizon; Xandros is in a good position to exploit these events. The software runs perfectly adequately on a past-it 500Mhz Celeron and allows users to stick with the applications they've grown accustomed to, while offering a route out of the *MS Office* stranglehold for users who don't need the more esoteric

features that suite offers. *StarOffice* is a good choice as an alternative in terms of support, but fails to match the finesse that the most recent version of *OpenOffice.org* can provide.

Where Xandros Business Desktop fares best is in the network stakes. We hooked into a variety of Windows and NFS networks with equal aplomb; add *Connector for Evolution* and you have a system that will offer facilities, performance, security and support on par with anything out there. The company should be able to make a compelling argument to businesses evaluating IT purchases in the wake of Microsoft's upgrade and licensing policies, and the hardware requirements of Windows XP.

It will be interesting to see how Xandros Business Desktop fares with increasing competition from Red Hat and Novell/SUSE in the corporate sector, but at present this is probably the easiest and most cost-effective route to getting Linux plus legacy applications onto a PC. We'd like to see COO 3.0 and *Connector for Evolution* integrated into the package as soon as possible to make it even more desirable to Windows users. The conservative selection of included apps seems designed for stability rather than the cutting-edge: no bad thing when you consider it may be a potential customer's first look at Linux.

The collection is well complemented by foolproof package management (thanks to *apt*) and the inclusion of the second CD which will save download time for small businesses still on dial-up. For more advanced admins, the ability to manage security updates and easily add Debian sources via the XN GUI will be a welcome addition. All in all, this is a very adept package which offers switchers a familiar working environment and a good set of productivity tools, while making it as easy as possible for admins to manage alongside an existing Windows network. **LXF**



StarOffice 7 is efficient, but surprisingly, its screen font rendering lags behind the smoothness characteristic of its sibling, *OpenOffice.org*.

LINUX FORMAT VERDICT

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

A very good introduction to Linux for business. The inclusion of the two *CrossOver* products makes this especially good value for money.

RATING **7/10**

●●●●●●●●●●

IDE/EDITOR

Visual SlickEdit 9.0

When is a text editor not a text editor? **Nick Veitch** reckons that would be when it's also a powerful Integrated Development Environment...

BUYER INFO

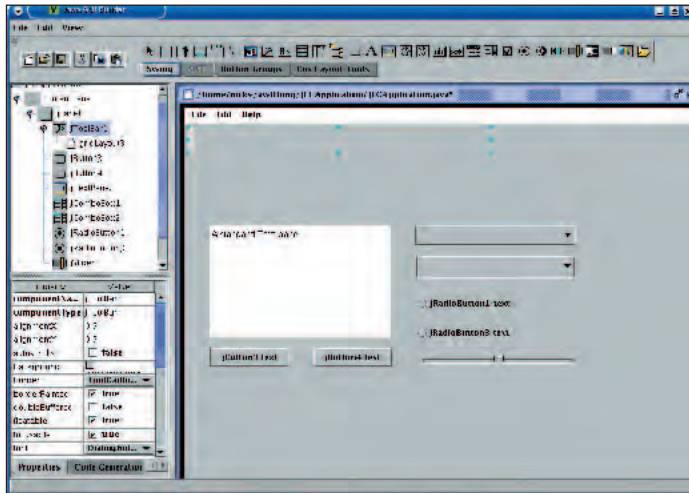
A full blown multi-language IDE. For specific development requirements you may wish to look at *KDevelop*, *Eclipse* or Borland's *JBuilder/C++ Builder* instead.

- **DEVELOPER** SlickEdit
- **PRICE** \$296/\$142 upgrade
- **WEB** www.slickedit.com/

Visual SlickEdit – or *VSlick* as many like to call it – has a long history on Linux. You may remember reviews of previous iterations in *Linux Format*. What started out as a powerful, programming-oriented text editor became a fully fledged IDE, with all the bells and whistles you would expect. With a large community to draw on, *VSlick* has incorporated many, many features and ideas from its wide user base over the years. Version 9 is no different in terms of its attitude, but greatly improved in terms of features.

Code refactoring

Apparently, *VSlick* is the first to include 'refactoring' for C++ code. Refactoring is the process by which you make a single change, which is then propagated throughout the entire code. Changing a variable name might seem like a simple operation, but it becomes more complex when you have to change it across 20 different source files, and make sure there are no conflicts. Imagine, then, trying another refactoring operation by hand – for example, moving a method from one



Build Java GUIs completely within *VSlick* – a real timesaver.

class to another, or converting a local variable to a class member object.

Refactoring should be a big time-saver, but we experienced difficulties getting it to work. Various errors always seemed to crop up, no matter how simple the operation we wanted to perform. Sometimes, the errors were not in our code at all, but apparently in the GCC include files! With any new feature, there are bound to be glitches, and we're sure a patch will be out soon to fix these teething problems.


Another time-saver is the new 'surround with' feature. Selecting a block of code, right-clicking, and selecting the appropriate option brings up a dialog containing mainly familiar constructs; these are things like if-else or switch blocks. Choosing one places the structure around the selected text

and places the cursor in a position to fill in appropriate details.

Java

VSlick is not just for editing C/C++ code. In fact, some of the major changes revolve around Java, the most notable of these being a GUI builder for AWT and Swing applications.

This is obviously a bit of a bonus. It is reasonably well laid out, and has all the features you might expect – tabbed toolbars for widgets, a property inspector where you can change crucial values and a hierarchical tree of your GUI design. It works, and seems to generate the required code without problems, but it is lacking in a few features. For example, once you've worked out the exact properties for number one in a series of ten buttons, you might reasonably expect that you could duplicate that object many times, and then simply shuffle them around, but there is no easy way to do this.



Code templates may save time if you remember to use them.

Unfortunately, there seems to be a few glitches in the Project wizard too. The wizard for a basic AWT application constructs a reasonable set of files for the job, but for some reason mis-references two of the classes created. Another easily solved error, but not what you might be expecting.

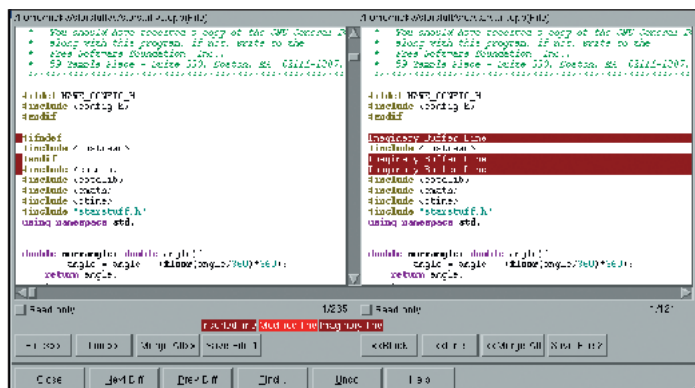
VSlick now includes a backup history for example, which means that outside of using version control, you can track changes to local files and view or restore to previous versions. The excellent *DiffZilla* difference viewer can be run separately and retain preferences for screen sizes. There is a new HTML-based help system (which needed reconfiguring on our system before it would work properly).

The update manager helpfully checks in with the SlickEdit website to see if there are any patches.

Competition

It's certainly a good IDE, but while once *VSlick* had little competition on the Linux platform, various Open Source projects are now mature enough to offer competition. For creating a KDE application, it's hard to beat *KDevelop* for instance. Similarly, Java developers are well supported by the *Eclipse* project (indeed, *SlickEdit* provides a re-worked version of this software as an editor plugin for *Eclipse*). Added to that, many of the tools sit happier on a modern Linux desktop – the *VSlick* GUI is functional, but a shade ugly.

VSlick is still a competent and professional tool that can handle a wide range of development projects. **LXF**



There are some tweaks to *DiffZilla*, the excellent difference viewer.

LINUX FORMAT VERDICT

FEATURES	9/10
PERFORMANCE	8/10
EASE OF USE	8/10
VALUE FOR MONEY	7/10

Configurable and feature-packed, but facing stiff competition from Open Source alternatives. New features help it keep ahead of the herd, but the gap's narrowing.

RATING **8/10**



LINUX DISTRIBUTION

Java Desktop System

Sun's attitude to the Linux phenomenon is confused to say the least. What effect will this ambivalence have on the company's flagship Linux project? **Andy Channelle** is becoming adept at doublethink...

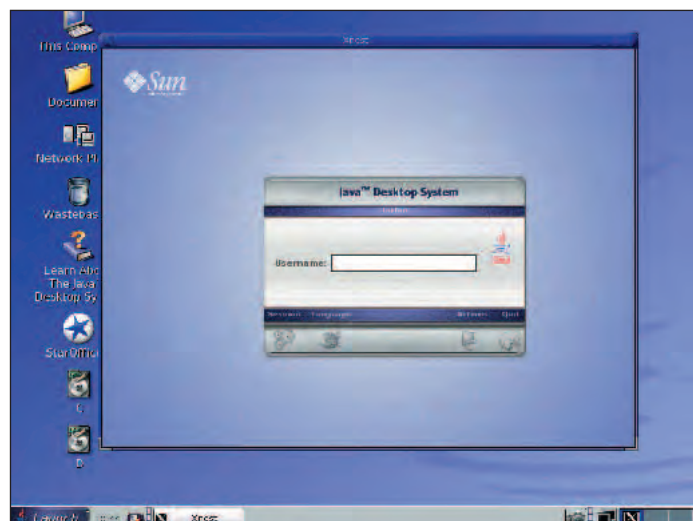
BUYER INFO

Business-centric desktop distribution aimed at Windows switchers. Also see: Xandros, SUSE Desktop

- **SUPPLIER** Sun Microsystems
- **PRICE** US\$100 per desktop per year (US\$50 until December 2 2004)
- **WEB** www.sun.com/

The presentation of Sun's new JDS release is faultless. The desktop is attractive and intuitive, the packaging is professional and reassuring, and the quality of the product is undoubted. This is because JDS is built on firm foundations – in fact we gave it's 'parent' SUSE Desktop Linux 1.0 a pretty good review last year – and features a wide range of very good software such as *StarOffice 7*, *Mozilla*, *Evolution* and *YaST* for configuration.

So why doesn't LXF like this quite as much as other desktops? Well, some pundits have taken issue with the licence that Sun has slapped on the software, suggesting it is the most restrictive software licence (including Sun's right to 'audit' our use of the software up to five years after the agreement ends!) since the dawn of time. As veterans of Sun's gargantuan licences, we're not surprised by this and, from a business perspective, it's no more onerous than the standard



The ability to start a new session within a nested windows is very cool, and very useful in day-to-day use.

MS fare – though it is a bit galling when discussing Linux. Others have suggested that, despite this being a 'new' release, it features quite a lot of 'old' software, including a 2.4 series kernel (2.4.19–4GB), GNOME 2.2 and *Mozilla 1.4*. Sun, I'm sure, would argue that its selection has the benefit of both maturity and stability and it would be hard to argue, especially with a product aimed at conservative buyers, for which this may be the first look at desktop Linux.

Our feeling is that it seems a little unfinished. Everything works as expected;

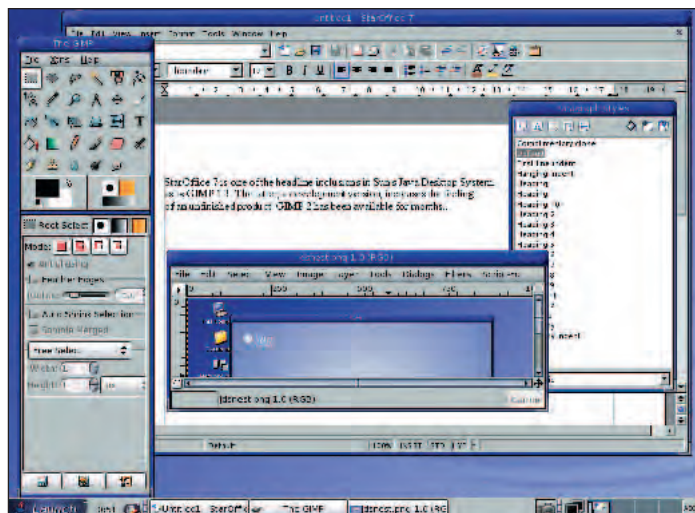
but scratch the surface, and the SUSE Desktop skeleton peeks through, echoing the lack of cohesion in Sun's Linux strategy. That said, in general use, JDS proved to be a good performer and the addition of the server-side *Sun Control Station* and the *Desktop Configuration Manager* makes light work of supporting a large number of systems, including applying security patches, installing new software and managing access policies. Integration with Sun's Calendar and Messaging servers (via *Ximian Evolution*) is similarly comprehensive.

Making room

One bizarre thing is, in perhaps 30 distro installations on a 3GB testing partition, JDS is the first ever to complain about cramped conditions. Running through the installation summary, the issue seemed to be caused by the default installation containing every localisation option available for GNOME, *Mozilla* and *StarOffice*, despite us clearly indicating language preferences at the start of the installation, and all the GNOME developer packages. After locating *YaST*'s software tools (sensibly buried away from prying users), we manually went through and removed all the unnecessary files to give the system a little headroom. Comprehensiveness can often be a virtue, but this is going a little too far!

Stack JDS 2.0 up against Windows XP/*MS Office* and you have a bargain. The idea of annual licensing may not be everyone's cup of tea, but if you have a 'soup-to-nuts' Sun operation, it makes a great deal of financial sense to fit the desktop in with existing infrastructure and the proprietary configuration tools are great for mass rollouts. Where it possibly falls down is in contrast to rival products from the likes of Novell/SUSE and Xandros. For example, the Xandros Desktop 2.0 Business Edition is not as visually refined as JDS, but it does include the same calibre of software, a more usable desktop, better *Samba* tools and a licensed edition of *CrossOver Office* (now including *CrossOver Plugin*) to boot, which might make it a more attractive proposition for those intending on keeping users working in *MS Office*.

These business-orientated products seem to have found a natural price point at around the US\$100 mark, so the decision comes down to other factors. If you're an existing Sun customer, the integration and configuration features and single-point support in JDS will probably win the day, but for those testing the waters, it really would be wise to test this alongside rival offerings from Novell and Xandros to work out which one suits the job at hand. So far, I feel, Xandros 2.0 BD is the one to beat, but Novell may have designs on the top spot this summer with its first unified SUSE/Ximian product. For individual users (admittedly not Sun's core market), there are better, more comprehensive introductions to Linux available at a keener price. **LXF**



JDS features a mix of the tried-and-tested (like *Mozilla 1.4*) for stability and the more unstable – allowing potential for future updates.

LINUX FORMAT VERDICT

FEATURES	7/10
PERFORMANCE	6/10
EASE OF USE	7/10
VALUE FOR MONEY	6/10

Good all-round distro that suffers in comparison to rivals in the sector. The management tools are great if you're already committed to Sun at the backend.

RATING 7/10



Books bonanza!

There are plenty of sources to guide you through all aspects of your Linux use, but some are much better than others. Our trawl includes the good and the bad!

We were inundated with a bumper crop of books this month, so rather drip-feeding you

our opinions of them over the next few issues, we decided to dedicate a whole eight pages to books this month! To top it off, we've also

included lists of our most recommended books in four key areas: *Programming*, *Learning Linux*, *Running a Server*, and *Security*.

ESSENTIAL BOOKSHELF MINI REVIEWS AT A GLANCE

Programming

- **PHP and MySQL** 33
- **Teach yourself C++** 33
- **The Secure Programming Cookbook for C++** 33
- **Learning Python** 33

Learning Linux

- **Linux in a Nutshell** 34
- **Unix Power Tools** 34
- **Red Hat Linux Bible** 34
- **Official Fedora Companion** 34

Running a Server

- **Official Samba 3 Reference Guide** 37
- **MySQL** 37
- **Managing Linux Systems with Webmin** 37
- **FreeBSD Unleashed** 37

Security

- **Practical Cryptography**.. 39
- **Maximum Security** 39
- **Practical Unix and Internet Security** 39
- **Beyond Fear** 39

REVIEWS AT A GLANCE

- **Essential Mathematics for Games & Interactive Applications** 32
- **RELAX NG** 33
- **Game Programming Gems 4** 33
- **PHP and MySQL for Dynamic Websites** 34
- **Guide to Elliptic Curve Cryptography** 35
- **PHP and MySQL Manual** . 35
- **The Definitive Guide to Samba 3** 36
- **C & Data Structures** 36
- **Essential PHP Tools** 37
- **Data Compression** 38
- **Definitive Guide to MySQL, 2nd Edition** 38
- **Game Programming Golden Rules** 39

Essential Maths for Games & Interactive Apps


All the numbers, functions, and graphs you could ever need...

BUYER INFO

- **AUTHOR** James M. Van Verth & Lars M. Bishop
- **PUBLISHER** Morgan Kaufmann
- **ISBN** 1-55860-863-X
- **PRICE** £36.99
- **PAGES** 676

Although we shan't get onto 3D programming in *Trout Wars* for some time, there's no harm you reading up about it in advance. When this book landed on our desks, we thought it was just another 'learn maths quick' book that either skimmed over the details or droned on and on about simultaneous equations.

Not so. Instead, this book is somewhat groundbreaking, for the primary reason that one of the two authors is a veteran games programmer who – most importantly – still works in the industry. The company he helped found, Red Storm Entertainment, has made some real chart-topping games, such as *Ghost Recon*, *Rainbow Six*, and *Rogue Spear*. The other author is a chief technology officer for NDL, that



produces a widely used game engine that has popped up in a selection of games our chums over at *PC Format* said were "mostly quite good, with some big names here and there." These two programmers together combine to form an elite maths team, which might lead you to believe the book is technically difficult. Thankfully, it's anything but: the text flows well and keeps your interest.

Though the writing style is to be admired, what's more important is the depth of coverage. We had some preconceptions when we picked up this book because we've read others in the genre in the past, most of which focus primarily on vectors and matrices. This

time breaks the mould entirely by covering all varieties of mathematics that relate to games creation – we were most surprised as we paged through at the sheer amount of thought that had been put into the contents list. Yes, vectors and matrices are included here (predictably), but so are the likes of affine transformations, rendering and picking, geometry and texturing, lighting and rasterisation, animation and curved surfaces, collision detection, and even rigid body dynamics.

We can imagine this book selling quite badly in bookshops because people will pick it up, leaf through the contents, and think, "there's no way

that can all be covered in meaningful depth in under 700 pages", and usually we'd agree. But the two keys to the authors' success are that:

- a** They don't fill up page after page with code (there's a CD for that) or padded-out text; and
- b** They know when to stop going into more depth about a topic.

The last point is quite important: this doesn't try to be the ultimate reference on everything mathematical. Instead, you get the basic use of a principle, plus a little more to help you stretch your limits, but then it refers you to other books that cover the area more – ideal, really. What's wrong with this book? We certainly couldn't find anything, and we challenge anyone to have a read, and prove us wrong.

LINUX FORMAT VERDICT

Comprehensive and to the point, this book is the best in its field without displaying any whiff of arrogance from the writing team.

RATING **10/10**



RELAX NG

Frankie say: you need to understand XML Schema languages...

BUYER INFO

- **PUBLISHER** O'Reilly
- **AUTHOR** Eric van der Vlist
- **ISBN** 0-596-00421-4
- **PRICE** £20.95
- **PAGES** 490

Believe that a good schema should not necessarily be a beautiful and aesthetic product of conceptualisation, but rather a working model by which data can be verified, analysed and used? Then you are probably already familiar with Relax NG. It's an XML schema that differs radically with the 'official' W3C XML schema – the latter is undeniably powerful, but has drawn a lot of criticism for being rather too complex.

The structure of the book leads us through the steps of understanding what Relax NG can actually do, then slowly builds from the creation of a simple schema to more advanced topics like modularisation and determinism.

Plenty of example material is scattered liberally throughout to illustrate key points, and all the topics



are presented in an easy-to-follow manner. Many of the examples relate to real-world situations, and it is this practical approach (though the author does drift into more conceptual topics), that makes the book so effective. Parts of it are even enjoyable to read!

The second half of the book is useful as a reference, as well as being a good learning source. Some 130 pages are devoted to the Relax NG element reference. This exhaustive work gives a brief description on element types, and includes an example for each.

Though it works up from first principles, the book may get a little too

technical fairly quickly if you are not already familiar with concepts like DTDs and schemas; but for those who have worked with XML or SGML, it should be easy-to-follow.

It's just as well this book is so good, because you won't yet find another volume available on this subject.

LINUX FORMAT VERDICT

Thorough and competent guide to a complex topic, written by an obvious expert in the field.

RATING **9/10**



Game Programming Gems 4

Is this the key to *Trout Wars* heaven?

BUYER INFO

- **AUTHOR** Andrew Kirmse
- **PUBLISHER** Charles River Media
- **ISBN** 1-58450-295-9
- **PRICE** £46.95
- **PAGES** 703

Last month we mentioned that *GPG 4* was coming soon and that nothing made us happier, and now it has arrived. This is no introductory text – if you're looking for easier topics, books one and two are probably your best bets. Book three was quite advanced, and this new release continues the trend.

It's not all hard work – there are about 20 or so 'easy introduction' essays that start from zero knowledge and solve complete problems by themselves. For example, *Advanced Wall Building for RTS Games* discusses how to program your AIs to defend themselves with wall structures and make the most of the impassable natural structures



around them. Although the average reader is probably only going to be interested in about a third of the content here – and likely to use perhaps only a quarter of it – the rest is available for interesting and insightful reading and is likely to give you some new ideas at the least. It's all cutting-edge stuff though, so for interest's sake alone it makes a very fascinating read, no matter what your games-programming skill level.

This is another great addition to the GPG library, and it continues to

push the skill level higher as you'd expect. Want to be on the cutting edge, or at least want to be able to work your way towards it? This is definitely the number one book to get.

LINUX FORMAT VERDICT

Comprehensive coverage of latest and greatest programming techniques for games, with something for everyone.

RATING **9/10**



PROGRAMMING

What are the top books to help you become the perfect programmer?

PHP and MySQL by Welling & Thompson (New Riders; ISBN 0-672-32525-X) is a sharply written book with a big following. Coverage is task-based, which means you learn as you implement parts of the overall projects. As a result, you can always see your learning in action immediately as you go. Despite a handful of minor problems, it's the easiest PHP book around to help get your skills up, and highly recommended.

Reviewed: LXF41, 7/10

Teach yourself C++ by Al Stevens (Wiley; ISBN 0-7645-2644-8)

is one of the better 'Learn C++' books out there; but, at just 700 pages, some parts of the language aren't explained in the depth you might expect them to be. For example, the discussion on post-incrementing versus pre-incrementing is covered in a mere page, and pointers is also quite terse. The only real drawback to the book is that the CD-ROM is mostly applicable to Windows users, which for a cross-platform language like C++ isn't good at all.

Reviewed: LXF43, 9/10

The Secure Programming Cookbook for C++ by Veiga & Messier (O'Reilly; ISBN 0-596-00394-3)

makes it quite clear from the beginning that your programs are insecure by nature. As it's a cookbook, each of the techniques discussed are backed-up immediately with code that implement it in C and C++. The topics covered are widely spread: such as validation, hashing, PKI, SSL, random number generation, and handling users trying to tamper with your binaries. Even if you don't plan on implementing any of the code, it's at least best to know what's available – this, above all, is what makes this book a must-buy.

Reviewed: LXF48, 10/10

Top Stuff Award

Learning Python by Lutz & Ascher (O'Reilly; ISBN 0596002815)

is one of the best short introductions to the language available, and serves as an excellent groundwork for those just starting out. Watch out for how to integrate C/C++ code, how to use Jython (Java for Python), and which IDEs can help the most. For beginners who are just starting out, this is almost certainly the best title available, and at just £24.95 for 600 pages, it's hard to go wrong.

Reviewed: LXF53, 9/10

PHP and MySQL for Dynamic websites

Half of the LAMP platform explained in a single book

BUYER INFO


- **AUTHOR** Larry Ullman
- **PUBLISHER** Peachpit Press
- **ISBN** 0-321-18648-6
- **PRICE** £18.99
- **PAGES** 570

This book sets out to be a practical guide to creating and running websites based on the ever-popular PHP and MySQL.

With a very practical, task-driven approach, it assumes little prior knowledge of either software, but certainly an understanding of HTML and how web sites work would be a great advantage.

Early chapters explain the basics of PHP and MySQL before going on to more advanced topics, and the book finishes with a series of hands-on projects to follow, such as building a CMS and an e-commerce site.

The introduction to PHP is fairly brief, but it is task-oriented and very practical. If you have any experience with scripting, it should be easy to pick



up the basics of PHP here (with a few visits to the excellent online documentation for the language at www.php.net as well, probably).

After the basics have been dealt with, the book returns to PHP for some real-world code dealing with common concepts. The section on cookies and sessions is very detailed, and includes plenty of helpful and

insightful advice. Typical of the book, each topic and example is worked through methodically, and if you follow the book from start to finish, you shouldn't find any need to skip about to work out what is actually going on.

Possibly for some reason we can't quite fathom, the author's code examples tend to be rather lengthy conditional blocks, which at times are

a little tricky to follow. It probably would have been better to modularise some of the lengthier examples by creating discrete functions, which would have greatly improved legibility in a lot of cases.

However, the security section is probably a little light on detail in some areas. Though all sorts of relevant topics are only briefly discussed (as they are encountered elsewhere, eg managing sessions), the chapter on security is a mere 30 pages; a little scant when later chapters are showing readers how to build e-commerce sites. Most of the security topics discussed in this chapter are about features and validation rather than making the code itself more robust.

LINUX FORMAT VERDICT

Certainly one of the better books on the topic, and great value.

RATING **8/10**



ESSENTIAL BOOKSHELF: LEARNING LINUX

What are the top books to get you started with Linux and help you hone your skills?

Linux in a Nutshell by Ellen Siever *et al* (O'Reilly; ISBN 0-596-00482-6), was described by us as "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," and is a classic collection of every major Linux command you're ever likely to need. This is one of the few books where *Sendmail* (a mail transfer agent), *sed* (stream editing), and *sdiff* (find differences between two files and merge interactively) share a double-page spread, with no convenient categories to help you find things easier. The advantage of the terse layout is, of course, that the book is kept relatively short, but we'd much rather have it spaced out neatly. Provided you know what you're looking for, though, you can find it all here, and that is what's most important. The best way to sum this book up is with the original reviewer's words: "If you don't lock your office, this will be the first thing that a techie colleague will steal!"
Reviewed: LXF46, 9/10

Unix Power Tools by Shelley Powers *et al* (O'Reilly; ISBN 0-5960-0330-7) is a set of individual hints and tips that cover a wide range of Unix topics as diverse as script writing, *Vi* and *Emacs*, printing, and networking. Split across 51 chapters, you might think the information inside was hard to find, but the inclusion of a 79-page index dispels all fears: it's all there for you to jump into easily whenever you need it.

Each of the tips, solutions, hacks, and workarounds presented have been discovered and refined through the years the authors have been using Unix, which means they are always practical and to the point – there's something in here for everyone, no matter how long you've been using Linux. Even here at *Linux Format*, we regularly pick up new inspiration from flicking through this book – sure, some of the ideas are very eclectic and require some thinking, but the vast majority are instantly helpful and can really make your Unix life easier.
Reviewed: LXF41, 10/10
Top Stuff Award

Red Hat Linux Bible by Christopher Negus (Wiley; ISBN 0-764-54333-4) covers both Red Hat's Enterprise Linux release and also its Fedora community release, which makes sound sense, as the two are so intertwined anyway. Various fundamental topics are covered, such as installation, configuration, and security, but they are covered in some surprising depth for what is hardly a big book – installation isn't just "put the CD in and follow the instructions". Instead you get treated to a full set of tutorials on installation from CD, FTP, NFS, HTTP, or Red Hat *Kickstart*, which is exactly what you need if you're an administrator in a complex computing environment. This is the big advantage of the book: you get to learn how to do things the way you want to do them, not how the author deemed it best to tell you. Despite the Fedora project moving swiftly forward with new releases, much of the content here is going to be valid for years yet, which makes this a very safe purchase.
Reviewed: LXF53, 10/10
Top Stuff Award

Official Fedora Companion by Nick Petreley (Red Hat Press/Wiley; ISBN 0-7645-5836-6). What other books do in 1000 pages, this does in 260. As this is the *official* guide from Red Hat Press, the content is more refined than you might find elsewhere – GNOME is, of course, covered in the most detail, but you'll also find quite a few pages devoted to the alternatives. Further examples of where the book breaks from the standard Red Hat party-line are easy to come by – look out for the guide on how to install MP3 support on Fedora, something that Red Hat continues to disable. Perhaps the main reason this book is so short is because the topics have been carefully chosen for home users: clustered web serving, distributed backups, and other complex scenarios aren't touched upon (quite rightly), which leaves this book feeling light and welcoming to less-experienced readers. Of course, it also means an incredibly low price tag – £13.99 – a computing resource well within everyone's reach.
Reviewed: LXF55, 10/10
Top Stuff Award

Guide to Elliptic Curve Cryptography


Number-crunching has never been so much fun...

BUYER INFO

- **AUTHOR** Hankerson *et al*
- **PUBLISHER** Springer
- **ISBN** 0-387-95273-X
- **PRICE** US\$60
- **PAGES** 311

Matt Groening's a genius: in the *Futurama* episode, 'Mars University, Fry asks Professor Farnsworth what he's teaching in the forthcoming semester, to which the reply is, "The same thing I teach every semester: the mathematics of quantum neutrino fields. I made up the title so that no student would dare take it." The title *Guide to Elliptic Curve Cryptography* is as close to this as it gets for books, and will likely turn many readers away from it before they even manage to open the cover. We steeled ourselves and gave it a thorough reading, and found it's not so bad as it sounds.

The important thing to understand about this book is that it starts off at a



beginner's level and works its way up to intermediate – it doesn't go much beyond that, largely thanks to its size, but it doesn't *need* to go much further because 'intermediate' in such a difficult field is usually hard enough for people! Although there is a fairly good ramp-up introduction for newcomers to cryptographic mathematics, it is very short and you'll need to get to grips properly in order to understand the later algorithms.

The further you get into the book the more it starts to sound like a mathematical text book, with such choice phrases as "let $\Psi : (P) \rightarrow (P)$ be

a group automorphism, where $P \in E(F_q)$ has order n ", but that's quite understandable – you can't discuss such a deeply mathematical topic without assuming a level of knowledge of the reader. It's possible that by adding another 100 pages of text the authors could have paced it much slower, explaining algorithms and algebraic notation to a broader, less knowledgeable reader-base as they went, but that would just have made the book more diluted for the intended primary audience – cryptographic professionals and mathematically inclined programmers.

To top the book off, there's a list of the 490 or so references scattered throughout the book so you can read more about specific topics if they interest you particularly, and they come from a wide selection of sources. What it lacks, however, is a concluding chapter that sums up what has been discussed, compares and contrasts the various methods, and perhaps also discusses the research that is currently underway into the field. These are minor points, though, and don't detract from the quality of the coverage given. Sure, the maths is a little hard for about 75 per cent of the text, but if your maths skills aren't that rusty, you should be able to pick up from anywhere and read only the parts of interest.

LINUX FORMAT VERDICT

Mathematically as hard as they come, but hits the target audience perfectly.

RATING **8/10**



PHP and MySQL Manual


Describes itself as a "comprehensive manual" that "explains how to use PHP to its full extent." Hmmm...

BUYER INFO

- **AUTHOR** Stobart *et al*
- **PUBLISHER** Springer
- **ISBN** 1-85233-747-8
- **PRICE** £39.50
- **PAGES** 589

There must be something in the air at this time of year that inspires people to publish books on PHP that are, quite frankly, not much good. It's been over a year now since the second edition of *PHP and MySQL Web Development* (one of the few PHP books worth its salt) was released, and since then, book after book has been churned out, essentially recycling material already available in the online PHP manual.

This book is no exception, and in many ways sinks lower than some of the other books we've seen previously. First, despite what the back cover might lead you to believe, there are substantial grounds to believe that the authors



have never tried PHP 5, let alone taken the time to write about it. The database coverage is strictly limited to *MySQL*, which is disappointing because PHP 5's new *SQLite* system is a work of art, and *PEAR::DB* continues to make strong advancements. It says a lot that if you look up *PEAR* in the six-page index at the back, you won't find anything; and indeed, we didn't spot any mentions of it at all in the main text.

The *MySQL* content itself weighs in at just a feather over 40 pages, and,

at its conclusion, mentions the `mysql_num_rows()` function: hardly rocket science. If you're looking for an explanation of normalisation, indexing, table types, transactions, hints and tips for optimisation – or any other quite commonplace activity – you simply won't find it here.

Aside from the serious lack of content, depth, and structure, there are still more annoyances. Side-stepping the whole glass houses and stones issue, we have to say that typing errors

seem to abound, even so far as the one-line code example for `imagecreatefromjpeg()` having the same line printed twice. The few pages that are free from typing tragedies are usually the ones with oversized, space-filling pictures. Take pages 16 and 17, for instance: on the left-hand page is a picture of the online PHP manual that takes up two-thirds of the page (with the helpful caption, "On-line PHP manual" [sic] and on the facing page is a picture of an offline PHP manual that also takes up two-thirds of the page (with the helpful caption "Windows help file"). Necessary? No. Helpful? No. Worth £39.50? We think not.

LINUX FORMAT VERDICT

Worse than *Core PHP Programming*, which is really saying something!

RATING **2/10**



The Definitive Guide to Samba 3


Cross-platform file-sharing made easy (and understandable).

BUYER INFO

■ **AUTHOR** Roderick W Smith
 ■ **PUBLISHER** APress
 ■ **ISBN** 1-590-59277-8
 ■ **PRICE** £31.50
 ■ **PAGES** 640

Anybody who is regularly involved in managing more than one computer will almost certainly be familiar with the general concepts of filesharing. If one of the computers involved runs on Windows, the overwhelming likelihood is that the type of filesharing being used is CIFS (Common Internet File System, Microsoft's name for SMB-based filesharing as used under Windows).

Samba is the weapon of choice for serving or attaching to SMB services on Linux and other Unix platforms, and for simple use-cases, all that is necessary is to install the software, do a minimal amount of configuration and off you go.



But *Samba* is – by necessity – a complex beast, and when you start worrying about larger networked filesystems, backups, authorisation, file-permissions, printers, nameservers and a whole host of other things. The aim of this book is to help you understand how *Samba* works, and get it to do what you want, solving problems along the way. As you might expect, the book begins with a description of what

Samba actually is, as well as the usual magic spells for downloading, compiling and installing the software. Only then do we get into the real hard information of configuration (which covers printers as well as file-sharing). With subsequent sections on backup techniques, talking to other protocols, troubleshooting and even scripting the *Samba* environment, this work is at least more comprehensive than most.

With plenty of suggestions and box-out tips too, it is eminently readable as well.

There is nothing particularly wrong with this book, but it is outclassed by the *Official Samba HowTo and Reference Guide*, John H. Terpstra and Jelmer R. Vernooij, Editors (Prentice Hall PTR, 2003, part of the Bruce Perens Open Source series. ISBN 0-13-145355-6). The latter is far more authoritative (being written by the developers), somewhat more extensive in scope, longer and available for around the same price. The only advantage Smith's *Definitive Guide* has is that it adopts a more friendly tone and hangs together slightly better as a complete work. Not to be disregarded.

LINUX FORMAT VERDICT

Outclassed by the official work, but very commendable nonetheless.

RATING **8/10**



C & Data Structures


If you thought C++'s STL was hard, wait till you see this...

BUYER INFO

■ **AUTHOR** P.S. Deshpande and O.G. Kakde
 ■ **PUBLISHER** Charles River Media
 ■ **ISBN** 1-58450-338-6
 ■ **PRICE** £40
 ■ **PAGES** 700

Although C++ is a complex language, it does benefit from having the Standard Template Library (STL). But what if you're using C, where the STL isn't available?

The actual coverage of the book is quite good, but perhaps not what you'd expect given the title – I don't know about you, but we all interpreted *C & Data Structures* as meaning "data structures written in C". What this book actually gives you is a 200-page generic introduction to the C language, followed by 500 pages of data structures written in C. While the latter part makes sense, and thankfully it's the larger part by a wide margin, the whole C introduction



section seems to be little more than space-filler. That's not to say it's a bad introduction – on the contrary, it's quite good, but it's redundant nonetheless.

The remainder is split into two sections: what the data types are and how they work, and problem-solving. It's hard to describe the first half as anything but basic, as large parts of it are code listings and output dumps. These input/output examples do help explain what's going on, but it makes a 10-minute job stretch over an hour.

Fortunately, this thins out once you get into linked lists and trees, where diagrams take up a lot of room. Having said that, as the implementations get more complex, the code stretches out even further – code listings of five pages or more are not uncommon.

The real treat in the book comes at the end, where problems with data types are discussed. As each of the data types presented are often used in real-world computing scenarios, you can learn a variety of techniques here –

linked-list garbage collection, Huffman coding, and even some good discussion on how to solve classic mathematical problems such as a knight's tour across a chessboard. Surprisingly, the A* pathfinding algorithm is discussed and solved also, which is something we shan't be getting onto in *Trout Wars* for about another six months!

This is worth buying for the problems section most of all. The actual data structures text is quite slow and watered down where you can probably find text of equal quality online. The opening part – discussing how C works, what variables are, etc – is quite pointless and will hopefully be excised in the second edition.

LINUX FORMAT VERDICT

The last 200-page block is a hidden gem buried under a pile of rambling that could easily be done without.

RATING **8/10**



ESSENTIAL BOOKSHELF: RUNNING A SERVER

What are the top books to get in-depth on key system administration tools and technologies?

Official Samba 3 Reference Guide from Terpstra et al (Prentice Hall, ISBN 0-131-45355-6) is a book of surprising depth. Surprising, that is, because it's a little hard to believe that anyone can get more than 700 pages out of *Samba*! But the collection of authors of this book have managed it, and their system is simple: cover the features, but, as you go, follow up each of the feature discussions with practical solutions to common problems along with hints and tips to make life easier. All the space has been used well, and we were thoroughly surprised at some of the ideas the authors came up with – using a *MySQL* server as a backend for account information, for example, is described as pretty crazy, but the authors explain how to do it anyway, just in case you want to try it yourself. All aspects of *Samba* use are detailed, from *Winbind* to domains and security, and from handling printing to the *Samba Web Administration Tool (SWAT)*. Although it's by no means the most entertaining read around, it's nothing if not comprehensive.
Reviewed: LXF53, 9/10

MySQL by Paul DuBois (New Riders; ISBN 0-7357-1212-3) stands out not only as an excellent title in the field of *MySQL* books, but as an excellent all-rounder even when compared to books of *any* topic. DuBois spends a lot of time talking and helping out on the *MySQL* mailing lists, which means he's not only intimately familiar with the day-to-day issues faced by users at all levels, but also very skilled at explaining solutions to people who are new to *MySQL*. This extensive knowledge is competently applied here, which means you'll find a full introduction to working with *MySQL* side-by-side with how to recover from disasters, how to work with the *MySQL C API* and more. As this is the second edition, the few mistakes that were present in the first release have been cleaned up and the overall text has been sharpened. Any programmer or sysadmin that has to work with *MySQL* should own this book and make sure they read it cover-to-cover at least once a year, and keep it on their desk the rest of the time – it really is that simple.
Reviewed: LXF41, 10/10
Top Stuff Award

Managing Linux Systems with Webmin by Jamie Cameron (Prentice Hall; ISBN 0-13-140882-8) makes managing your servers remotely much easier. Part of the reason for this is that Jamie Cameron is actually also the author of the *Webmin* software itself, which puts him in the position of knowing it inside-out before he even starts. Of course, this isn't always an advantage, as not every highly skilled software programmer can naturally write good English, but this is a trip-up that Cameron deftly avoids – every official *Webmin* module is covered here in great detail, and it's not at all dull reading. Backed up by copious amounts of screenshots, tables, and diagrams, it's pretty easy to follow the various step-by-step guides, and soon you'll find yourself quickly adding subnets to your DNS configuration, setting up proxy authentication for *Squid*, etc. Although you could decide to read the book straight through, it's much better if it is used as a dip-in reference. If you're considering *Webmin* for more than just home use, this book will save you a lot of trouble.
Reviewed: LXF46, 9/10

FreeBSD Unleashed by Urban and Tiemann (Sams; ISBN 0-672-32456-3) is the perfect book to keep to hand if you want to experiment in the world of Unix outside of Linux. Thanks to the similarities shared between all Unix systems, large chunks of your Linux knowledge will transfer smoothly, but some won't. That's where this book comes in: wherever there's a peculiarity to FreeBSD, it's covered here. Sure it's big, but the space is used wisely. Most helpful for us is the good list of comparisons between where Linux and FreeBSD are applicable, and there is of course a large software crossover between the two systems: KDE, GCC, etc, are the same wherever you go. The 1000 pages are filled with some out-of-the-way topics also, which is great to see. The section on shell scripting particularly stands out as something that's missing in other books, and you'll even find a quick Perl intro here also. Even if you're new to Linux, this book makes things quite easy, and a lot of work has been put into emphasising and explaining the similarities between Linux and BSD.
Reviewed: LXF44, 10/10
Top Stuff Award

Essential PHP Tools

That's Modules, Extensions and Accelerators, and not *Vi* and *Emacs*, whatever the hardcore may say!


BUYER INFO

- **AUTHOR** David Sklar
- **PUBLISHER** Apress
- **ISBN** 1-59059-280-8
- **PRICE** £22
- **PAGES** 347

While unquestionably a powerful web scripting language, PHP really comes into its own when

you begin connecting your own sites, databases and programs up to advanced third-party modules. *Essential PHP Tools* focuses on a handful of the best, broken down into five categories: databases, HTML, XML, Networking and the catch-all Debugging, Caching and Optimizing.

Of course, this presents us with the first catch: if the modules in those categories don't interest you, you aren't going to get much out of this book. It makes no attempt to teach advanced



PHP programming or widen the scope to other modules that you may encounter, instead focusing specifically on the ones that it recommends.

Luckily, this doesn't extend to the individual modules on offer. Taking *Smarty* as an example, we begin with a look at what it is – a templating engine used by an increasing number of scripts and content-management systems – how it works and why you should use it, before hitting a series of basic tutorials to explain the core

concepts. The book promptly whips through the commands, modifiers and tags that you'll need to know, before diving into a more advanced tutorial on how to go about implementing it on your own site. By the end, you have a fully working micro-application, in this case an online ice-cream store, and code that can be easily adapted for your actual website.

Obviously, the complexity varies according to the module being used – sending email and processing email

forms being on a slightly different level to developing web services with SOAP, but the tutorials are all extremely easy to read. While there'll always be one module that you really wish had been included, and one particular task that you'd love to have explained in additional detail, *Essential PHP Tools* provides a decent spread of features for any intermediate level PHP programmer. True, it's often the kind of advice you'd be able to work out for yourself given time, manuals or a convenient online forum, but having it neatly packaged up in one place is unquestionably an advantage.

LINUX FORMAT VERDICT

Have a scan through the contents to make sure it really covers a couple of topics you'd consider 'essential'. If so, it would be worth considering.

RATING **7/10**



Data Compression 3rd Edition

Yes, it's £54, but it's one of the few books that's actually worth it...




BUYER INFO

- **AUTHOR** David Salomon
- **PUBLISHER** Springer
- **ISBN** 0-387-40697-2
- **PRICE** £54
- **PAGES** 700

Back in LXF53, we reviewed David Salomon's previous book: *A Guide to Data Compression Methods* (Springer; ISBN 0-387-95260-8) and, although it scored 8/10, it paled in comparison to his second edition of *Data Compression*. This book is the third edition in the line, and is really what we've all been waiting for.

Although we could burble on for a while about Salomon's depth of knowledge on the topic, it's easier if you just understand that he's the leading author on compression! This book is a full update of the second edition, bringing with it lots of new content about new techniques, but also updated content on older techniques. The PNG image file format now gets full exposure including a look



at how PNG interlacing works. There's also a selection of new compression methods discussed, including PPM, SCSU, BOCU-1, MLP, Xmill (for XML), and others.

Despite now being in its third edition, substantial changes to the main bulk of the text are still occurring. As it's so data-heavy, it's unsurprising that previous editions had quite a few errors in that were reported on the book's website; these are now, of course, all corrected or updated in the paper edition. However, some things have yet to change, which is a little annoying: the discussion of MPEG-4, for example,

remains stubbornly non-technical. Part of the problem with the format is that it's designed to do a lot more than just compress video, and even when it does the compression it uses a variety of algorithms depending on the quality of video that's needed. However, it would be good to see at least one of the algorithms broken down!

Some of the best text in the book – in our opinion – is the section on JPEG 2000. Although the standard has really yet to take hold, it's only a matter of time, and it's given a thorough going over here. As it uses wavelet compression (also something that's yet

to achieve widespread popularity) reading about JPEG 2000 teaches you a lot more than how to just compress pictures, which explains why the author went into so much detail.

In our view, the newly added material alone isn't really enough to justify the purchase of this book if you already own the second edition. If you own the first edition, then you should definitely 'upgrade'. If you don't own it at all, you might find the mathematics involved a little too hard for you – borrow someone else's copy before splashing out on your own. What is worth the upgrade for second-edition readers are the corrections: if you work with data compression regularly this book is – and will probably always remain – a must-buy.

LINUX FORMAT VERDICT

Comprehensive coverage on everything but MPEG-4, but we like it anyway – bravo, David Salomon!

RATING **9/10**



Definitive Guide to MySQL 2nd Edition


"Definitive" and "MySQL" always perks our ears up...

BUYER INFO

- **AUTHOR** Michael Kofler
- **PUBLISHER** Apress
- **ISBN** 1-59059-144-5
- **PRICE** £35.50
- **PAGES** 802

Just a year or so ago, LXF didn't review many Apress books because the company just seemed to be churning them out in a manner that made us think it favoured quantity over quality. This has changed recently, and the quality of its titles is improving; and, at first glance, this appears to be great: the layout is attractive, the type clear, and there's a good mix of pictures and diagrams to help break up the text.

But you don't read a book like this to look at the pictures: at 802 pages, how much depth does this book go into? The answer is "quite a lot, but not as much as Paul DuBois' book", which is a good



thing and a bad thing. Although the two books are almost of an equal length, they choose their specialties differently. Kofler doesn't go into much depth on optimisation techniques to get your MySQL server to eke out every drop of performance, but on the flip-side, there's an excellent chapter called *SQL Recipes* that follows on directly after the chapter on elementary SQL.

Similarly, while this title doesn't go into that much depth on all the different

parameters that can be passed to the various MySQL programs, it does cover language bindings for the MySQL API in PHP, Perl, Java, C and C++, Visual Basic, and C#. Discussion on PHP inevitably forms the lion's share of the language content, and it's to the point and well-written. Though this book was originally written in German, you would never be able to tell from the quality of the text – it keeps your interest well, showing that the translator did a sterling job.

Although the coverage as printed is good enough, there are some bits we'd like to see improved in a third edition. First, speed is one of the key reasons why people use MySQL, so a chapter on optimisation would be very welcome. Second, the chapter on MySQL 4.1 is tacked on to the end like an afterthought, and could be much improved. MySQL 5.0 was announced quite some time ago, and code has also been available for a while, and yet for some reason it remains unmentioned – this is the sort of thing users would like to know about so they can plan an upgrade strategy.

LINUX FORMAT VERDICT

Not quite definitive, but it makes an excellent effort and is highly recommended.

RATING **9/10**




Game Programming Golden Rules

Silence is golden, but so are games programming rules too, it seems!

BUYER INFO

- **AUTHOR** Martin Brownlow
- **PUBLISHER** Charles River Media
- **PRICE** £33.50
- **ISBN** 1-58450-306-8
- **PAGES** 318

What are rules? If a rule is there to be obeyed, sometimes bent, and on rare occasions broken when there's an over-riding justification to do so, a *golden* rule must therefore be intended to be obeyed all the time. These inviolate commands are there to make our lives easier and help us avoid embarrassing fates: for example, when in a queue with people you don't know, you might either say: (Europe) "It looks like we might have a spot of rain in the afternoon," or (America) "How 'bout them Dodgers?" So, a book of games programming golden rules must surely be packed with invaluable knowledge and years of learning just waiting to jump out and surprise you.



Instead, what we have here is a rag-tag collection of misfit 'hints', of which the majority are flat-out obvious. Any programmer worth their salt – and we should keep in mind here that games programmers are usually amongst the most skilled programmers around – will already know exactly how the C/C++ pre-processor works, what macros are, etc. Similarly, the book goes to great lengths to explain what the 'const',

'register', and 'volatile' modifiers do: something programmers learn about in their second class on the subject!

In comparison to books like *Game Programming Gems* (from the same publisher), this title pales almost into non-existence. The tips on optimisation are frankly worthless; but if nothing else, at least they come right at the beginning of the book so that people flicking through at the bookshop will

quickly figure out the book's true value! On the CD you'll find full copies of the code used in the book – don't get your anticipation up here – and also, joy of joys, a full copy of the *Zlib* source code. As if you wouldn't just download it from the web if you wanted it.

Where the book has a chance to shine, it is cruelly shot down by the author. For example, when discussing handling game assets such as pictures, the implementation is done using Microsoft's proprietary *DirectX* system "for ease of use". There really is nothing here – *nothing* – that hasn't been done better elsewhere, which makes this book wholly redundant and deservedly consigned to the bin. **LXF**

LINUX FORMAT VERDICT

Golden what? We have completely forgotten this one already.

RATING **1/10**



ESSENTIAL BOOKSHELF: SECURITY

What are the top books to help understand security issues?

Practical Cryptography by Ferguson & Schneier (Wiley; ISBN: 0-471-22894-X) is one book we won't forget. Cryptography is never an easy topic, and we'd be lying if it we said it was made into an easy topic in this book. However, these authors are what can only be termed *cryptosmiths* – masters of the cryptography trade with years of combined experience. More importantly, they share all that experience with us in this book! Together these two have worked on both the Blowfish and Twofish algorithms, the latter of which was a strong contender for the Advanced Encryption Standard competition. How encryption works – from the basic nuts and bolts of the mathematics behind it upwards – is covered here, along with practical advice on how to implement it in your own situations. Yes, the topic is hard and nothing can be done about it, but the authors write with so much ease and so much understanding of the topic that it's actually easy to just sit back and soak up their knowledge. **Reviewed: LXF45, 10/10**
Top Stuff Award

Maximum Security by Anonymous *et al* (Sams; ISBN: 0-672-32459-8) is a heavyweight tome that covers a selection of security topics that often get overlooked. We have to give respect to any book with a chapter called *Practical WEP Cracking*, but the goal is to teach security by explaining how crackers get in and what they want to do. The introduction covers most of the non-technical data, such as designing a security response policy and coping with the aftermath of an attack, but the book swiftly moves on to how TCP spoofing works, firewalls, mailbombing, viruses, and DOS attacks. There are already quite a few books on the market that supposedly teach you how to crack into other computers, but they generally lose the plot and you can tell the authors almost want to break into 31337-speak. This one stays firm in its entirety, which means you come away with some powerful new learning about how to protect your network without needing to sit through week-long courses on the topic. **Reviewed: LXF42, 10/10**
Top Stuff Award

Practical Unix and Internet Security by Simson Garfinkel *et al* (O'Reilly; ISBN 0-672-32459-8) is now firmly established as a classic. Of course, the information presented in here is applicable to any Unix clone or family member, so it's wholly applicable to Linux. The difference between this book and something like *Maximum Security* is that this is aimed solely at intermediate to experienced users, which means there are no explanations here of simple topics. Furthermore, it is of course filtered for Unix-only content, so the 1000-odd pages of content are equal to 2000 pages of Unix/Windows content. Much of the content is philosophical in nature: always choose good passwords, try to close as many ports as possible, disable unnecessary services, etc, but it's all explained in a relaxed tone that helps the guidelines sink in. As Charlie Stross said when reviewing a previous edition of the book: "If you know nothing about Linux security, and only have time for one book, you should start with Practical Unix and Internet Security." **Reviewed: LXF42, 9/10**

Beyond Fear by Bruce Schneier (Springer; ISBN 0-387-02620-7) was written fresh from the success of *Practical Cryptography*, and is a less-technical discussion on the problems that surround computer security in everyday life. Questions that might sound obvious end up being complex, but Schneier guides us safely through it – is shopping online safe? Will computerised elections make voting more accurate? Schneier's answers usually sound a little crazy at first (hint: whatever you think the answer is, you're probably wrong), but the bulk of the book is spent explaining these views and showing why we need to put a lot more thought into these topics. Each topic is backed-up with copious amounts of source material from elsewhere so you can double-check his references; but there's little need to: Schneier is a master of security, and here he's in his element. This is a must for every serious computer user; but, more importantly, is an absolute must for every politician also. **Reviewed: LXF49, 10/10**
Top Stuff Award

Roundup »

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

OUR SELECTION AT A GLANCE

- Inkscape
- Karbon14
- Sodipodi
- OpenOffice.org Draw
- Sketch/Skencil

Illustration software

Illustration software for Linux has had a slow start. Now, options are expanding: can an Open Source project provide the high-end features of *Illustrator* or *Freehand* without the development muscle of Adobe or Macromedia? Andy Channelle sharpens his pencil...



Many illustration software packages appear to inspire a loyalty that goes beyond the features of the latest release. Professional and casual designers alike tend to fall into a specific method of working that fits with, often, the first tool they pick up. Witness, for instance, the failure of Corel to break the *Freehand/Illustrator* duopoly in the Mac market, despite leading the pack in terms of features and usability for much of the 1990s. *CorelDraw* was, of course, a Windows-native package, and was thus regarded as a toy by snifty Apple users. So what hope does an Open Source project such as *Karbon14* or *Inkscape* have in such hostile territory? Well, the answer is – at the moment – very little. Just as *The GIMP* suffers in comparison to *Photoshop*; *Sketch*, *Sodipodi et al* will grow in popularity on Linux but will probably struggle to win over diehard users of other applications of just their merits alone. All this would

matter, of course, if any of these applications were destined to live or die on their commercial prospects. Fortunately, they're all Open Source, and so are not tied to the precarious fortunes of the bottom-line. The Open Source method also means that if one application gets a killer feature, chances are it can be integrated fairly painlessly into its rivals; this is especially true of projects based on the same codebase: such as *Inkscape* and *Sodipodi*.

As might be expected, there are applications in our quintet developed with specific desktop environments in mind, such as *Karbon14* (KDE) and *Sodipodi* (GNOME), and these benefit from a look and feel that is consistent with the environment and also integration with other tools associated with their respective desktops. Other applications such as *OpenOffice.org Draw* aim to offer not just consistency across Linux systems, but also across platforms, which will have benefits in

heterogeneous environments. All, however, work on the same principle: the user builds up an image with a range of primitives – squares, circles, curves, etc – colours and lines, which remain editable at every point of the creation process. To this, the applications add various ways of creating colour graduations, transparency and patterns to build up more sophisticated imagery.

There is a difference of ambition in many of these projects. For instance, *OpenOffice.org Draw* is intended to be used within the confines of a larger office suite and therefore won't have the same audience in mind as the developers of *Sodipodi* who may harbour the lofty ambition of creating a rival to the commercial applications available. We would expect, for instance, *OOo Draw* to make a strong showing in the area of text design and desktop publishing (DTP) tools, while those that set their sights on being a

more specialised illustration tool should have a more comprehensive featureset and, ideally, a user interface that is comfortable for experienced users of other applications.

Despite these differences, the illustration applications on test in this issue's *Roundup* have many similarities: there are only so many ways you can do a job after all, and all rely on a subset of XML or SVG as their native file formats. This means, for instance, that working across applications and creating vector graphics suitable for use on the Internet (see boxes, pages 45 and 46), which are both scalable and bandwidth-friendly, should be – but is not always – simple. And while support for a recognised standard is important, it is also vital that these applications play nicely in the real world, that means at least being able to import or export the formats that have become *de facto* standards in the illustration/graphics world.

Inkscape

A fork from *Sodipodi* with a focus on usability.

■ VERSION 0.38.x ■ WEB www.inkscape.org/

Inkscape gains prominence over its progenitor by dint of its superior interface and the pace of development that the coding team has established. This is especially remarkable when you consider the group has only been active on the project for seven months. The latest version is available in a number of formats, including RPMs, debs and a Windows executable, from the project's homepage. It installed without issue on everything we tried.

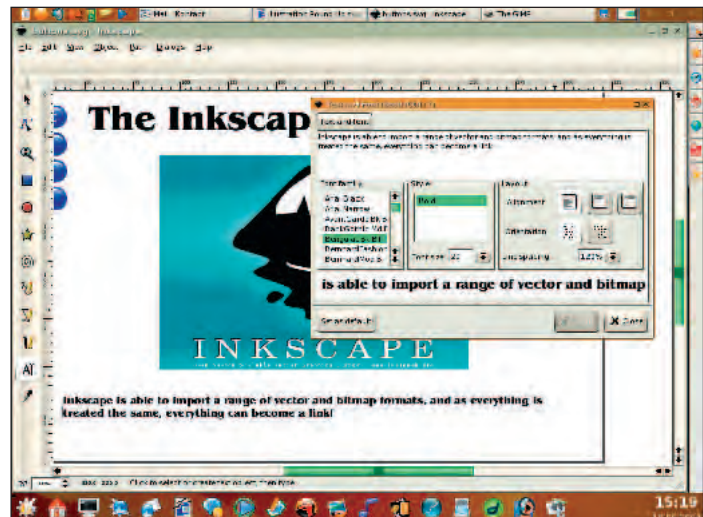
In contrast to *Sodipodi's* many-windows approach, *Inkscape* ranges a selection of tools down the left side of the screen (in the traditional fashion) with a context-sensitive toolbar running across the top and an ever present view of the 'page'. This is augmented by a selection of extra toolboxes covering colour selection, stroke width, etc, which can be launched or dismissed as needed. This approach makes for a more generous workspace and less dialogue box hunting, while also providing finer control when necessary.

Primitive shapes available include the usual stars, spirals, squares, circles and bezier curves, and all of the tools have a selection of options available on the context sensitive top bar. The spiral tool, for instance, includes options to change the number of turns, the 'divergence' of the outer parts of the spiral (how far apart the turns are) and the tightness of the

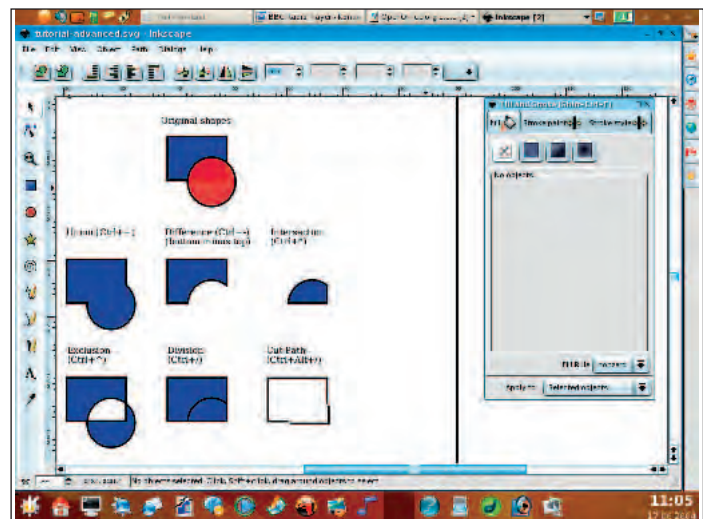
'inner radius' (the first turn) of the image; while the 'star' tool has options for the number of corners, the sharpness of the points and whether the shape is a polygon. The images are also interactively adjustable, so clicking the inner part of the spiral allows us to 'roll it up' from the inside, while grabbing a point of the star allows us to change the sharpness of the points and move them around to create a ton of cool shapes. Squares, meanwhile, have two handles which change the diameter of a curved corner, which makes creating website/application buttons in the ubiquitous Aqua style quite simple with the addition of a little transparency. Crucially, *Inkscape* has the most comprehensive set of 'path' tools for cutting, combining and generally mangling shapes, and this does lift it above the rest of these apps quite significantly.

We like the ability to add gradients, complete with transparency, to both fills and strokes (lines), though a method for naming gradients and a more (forgive the pun) transparent way of managing them would be very welcome.

If you are creating buttons for use on an SVG project, *Inkscape* makes it easy to add a hyperlink to any graphical element. It would be quite easy, though not yet very practical, to create a whole website in this way if you so desired.



The text tool is quite well featured, but we really miss a text-to-path option – hopefully it's on the list of future features.



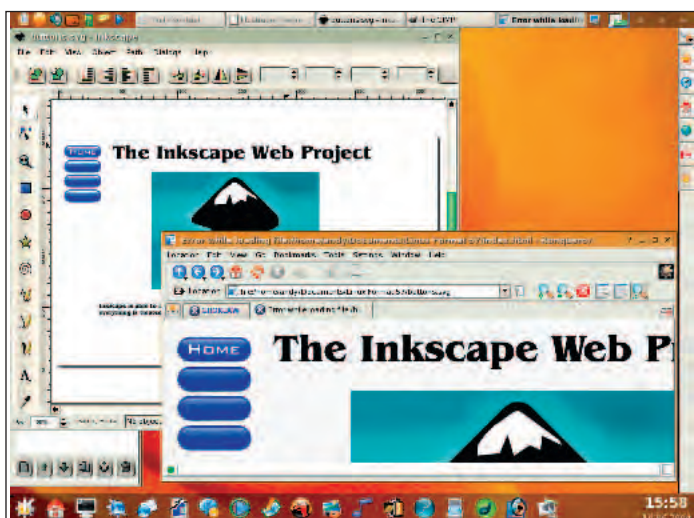
Boolean operators lift *Inkscape* above the *hoi polloi* – a level of sophistication that is an example to many other Open Source apps.

In terms of text, *Inkscape* supports both horizontal and vertical text, and words remain editable regardless of how you twist and turn them around. One thing that we did miss though was the ability to automatically fit text to a curve. The 'Tips and Tricks' help file contains a well thought out workaround, but it involves turning the text into non-editable curves, which makes going back a bit of a chore.

One feature that we really liked (and one we'd like to see expanded upon in the future) is the calligraphic pen, which adds a thick, directional nib to freehand strokes. It's really quite effective, and with the addition of a selection of nibs, would make this a pretty smart 'natural media' drawing application, or Corel-like image hose.

The future of *Inkscape* looks bright, especially if the planned cooperation with *Scribus* and *The GIMP* comes to

fruition, which would finally realise the concept of a full-spec graphic suite to Linux. Import and export options are adequate, with bitmaps even inserted into SVG files correctly. The crucial omissions in this early version are support for PDF and EPS formats, though getting around this is not too onerous with a little help from *Scribus*.



Inkscape's SVG output can be viewed directly in *Konqueror*.

LINUX FORMAT VERDICT

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

A young, but very promising application that has the potential to outclass the other apps reviewed here. Definitely one to watch for future improvements.

RATING **8/10**



Karbon14

Kontour's replacement – a KDE killer app?

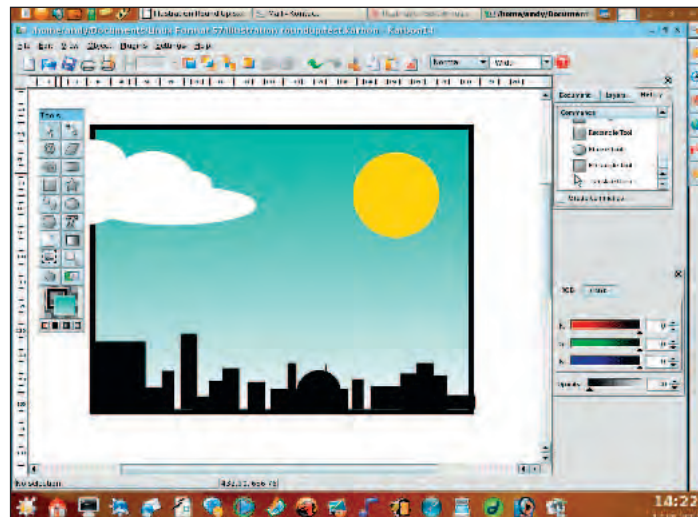
■ **VERSION** 1.3.1 ■ **WEB** www.koffice.org/karbon/

Karbon14, as its initial letter might suggest, is a part of KDE and more specifically the illustration element of *KOffice*, where it replaced *Kontour*. It has a good selection of tools available and integrates quite tightly with *KOffice*. It is possible, for instance, to open an instance of *Karbon14* from any other application in the suite within the page to create or edit images with the full range of illustration tools. Double-clicking on the embedded image in the page relaunches the *Karbon* tools so the picture can be edited *in situ*. This is a killer feature and, in most cases, it worked as expected. However, on the couple of occasions we managed to crash *Karbon14* during embedded editing, it took down the host application as well which was a little frustrating. It would also be nice if this ability was reciprocated, but *Karbon14* doesn't yet act as a host to other *KOffice* applications.

The application offers a good selection of illustration tools, though the plain 'pen' tool is rather obscure

called the Polygon tool, as it handles both bezier curves and filled, irregular polygons. Once shapes are created, they become editable using the selection of 'node' tools. These aren't as comprehensive as the options available in the other applications; lacking, for instance, tools to add new nodes, break them apart or join them together. We also found it quite frustrating to use the node tool as it's incredibly fussy about how the user can select nodes and, if we're on the subject, the objects themselves.

We encountered quite a few stability issues in general use. A number of processes could reliably kill the application stone dead. Among these were: using the interactive zoom tool, though zooming via the View menu was OK; importing SVG files; and occasionally importing bitmaps. On the other hand there are some features that, though in their infancy, will make *Karbon14* a really useful application in the future. We were excited by the prospect of outputting



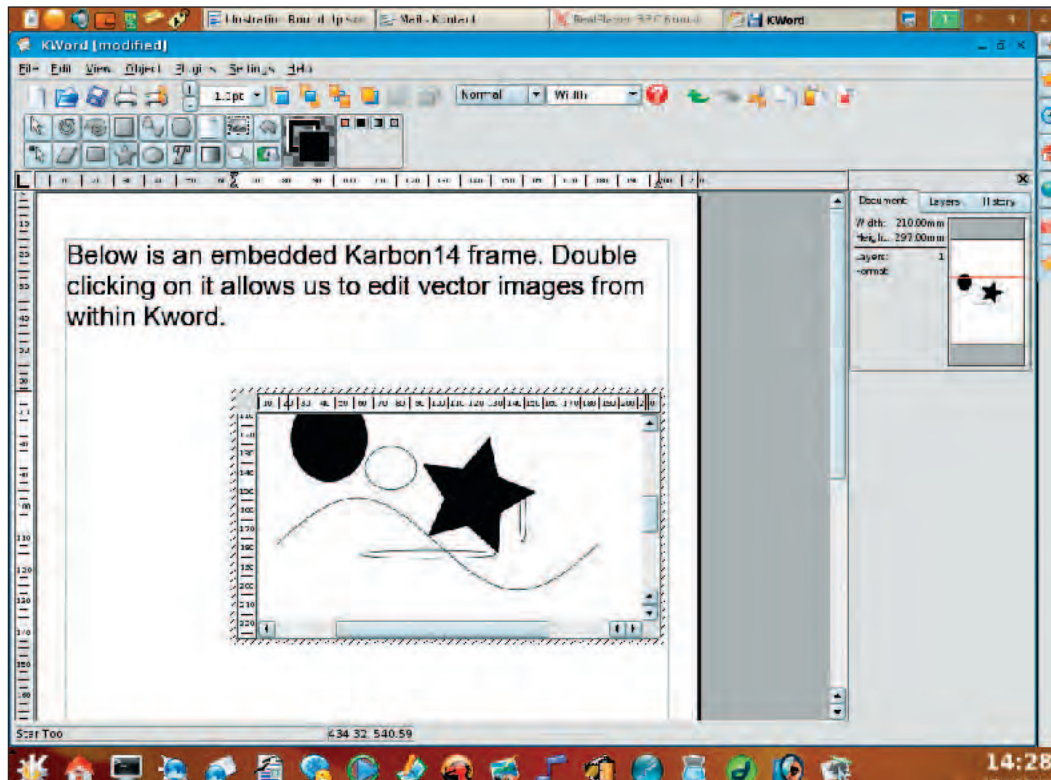
Creating this very simple scene was actually quite a difficult task, thanks to the rather idiosyncratic selection tool.

native *GIMP* bitmaps (hoping perhaps the layer data would remain intact), but *The GIMP* subsequently refused to load any of our files, claiming the format was corrupt.

It's not all bad. We particularly liked the ability to shift the different parts of the user interface around. For example, on first launching the application, the main drawing tools are housed in a secondary toolbar at the top of the window. Double-clicking anywhere in the toolbar turns it into a floating box, which can be positioned anywhere on

the desktop, where it remains 'on top' of the window stack. Double-click again and it flies back to the original location. The same is true of the option boxes which, by default, are docked to the right hand side of the window. The history tool concept is also excellent and should be implemented in every computer application. This allows the user to shift back to the moment before a bad decision was made and undo everything that followed; a far better solution than blindly hitting **Ctrl+Z** and hoping the application supports an adequate number of undos. The history can consume quite a bit of memory, so if things become tight, it makes sense to hit the 'Purge' option in the Edit menu to clear out the buffer.

Karbon14 is an application that is full of good ideas, and with some refinements and effort expended on improving stability, it could be a real contender for the top spot. If you want or need the nicely handled integration with *KOffice* and KDE, there's no other game in town. Just remember to save your work often!



Embedded editing within KWord takes a little getting used to, but shows a great deal of promise.

LINUX FORMAT VERDICT

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

A nice app that's sadly let down by some annoying bugs in its present incarnation. Future development could make turn this contender into a heavyweight champion!

RATING **6/10**



OpenOffice.org Draw

Illustration component of the all-conquering productivity suite.

■ **VERSION** 1.1.1 ■ **WEB** www.openoffice.org/

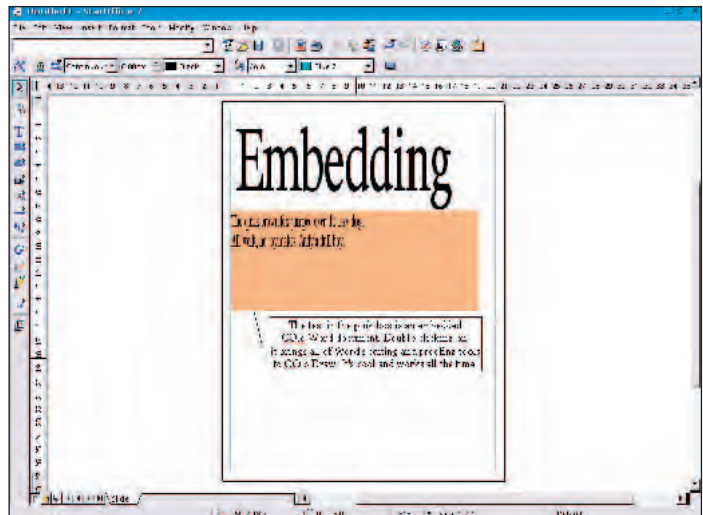
Jack of all trade applications, it could be argued, tend to sacrifice detail in pursuit of becoming all things to all people. Happily this isn't the case with *OpenOffice.org Draw* (henceforth called simply *Draw*) which manages to pack in a wide range of tools suitable for illustration, diagram creation and presentation, while retaining a clean user interface.

In fact, as an all-rounder, there's nothing in the Linux world to touch *Draw*. It has the full complement of drawing tools, curves, stars, spirals etc., but also includes a useful selection of label/diagram tools and text-based 'Callouts' and 'Pullquotes', so it becomes useful for creating flow charts, network diagrams etc. Gradient and transparency tools are both comprehensive and easy to manage and, for those in a hurry, a decent selection of preset fills save the effort of creating your own. *Draw* is also the only application on test here with an intuitive method of saving preset fills.

The user interface – as one might expect for a component of a suite that

started life as a text-editing program with a few bells and whistles – reflects more of the 'word processor' than 'illustrator' paradigm, and this may be seen as a bonus to the general office user who may only be called upon to add a background and some fancy fontwork to a presentation. Like *Inkscape*, the main tools are displayed down the left-hand side, though the context bar at the top get less use, as *Draw* relies more on button flyouts. These are fully moveable (which can be a blessing or a curse depending on how you work) and contain all the options for the given tool.

Integration is undoubtedly *OpenOffice.org*'s strongest suit. To add a text element to a *Draw* illustration, we can either use the application's own basic tools or, for more comprehensive options, we can embed a fully resizable *Writer* frame into the document to bring all of that applications text editing tools to the page. This facility also works the other way round, and indeed between any of the applications in the suite.



Application integration is a desirable goal, whatever your OS, and it's *OOo*'s *Draw*'s strongest point. One app to rule them all!

Other highlights include the 3D primitives, the lovely shape distortion options and the often overlooked line and line cap options.

On the import/export front, *Draw* works with most of the standard formats, although – and this is a little

bizarre – there is no facility for importing bog-standard SVG files. It can, however, export them along with PDF, EPS, SWF (Flash) and a host of bitmap formats.

Though part of the 1.1.1 release, it is not bug-free though, and the most annoying glitch affects one of the most vital tools. To see it in action, select the 'Curve' tool, make your initial mark on the page and drag the cursor across the page. Now click and hold to create the second point and then drag the cursor down. In *Draw 1.1.1* (and also in *OOo*'s proprietary cousin, *StarOffice 7*) a third point is automatically added to the curve in the top left corner of the page. This can be really annoying when all you're trying to create is an arc, and it prevents *Draw* from really shining in this area.



OpenOffice.org Draw is the only application in this Roundup with its own (rudimentary) text-to-path tools.

LINUX FORMAT VERDICT

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

The best general-purpose drawing package available for Linux today. As one part of the class-leading office suite for Linux, the future looks very bright indeed...

RATING **8/10**



Sodipodi

A package that follows the design of early *GIMP*.

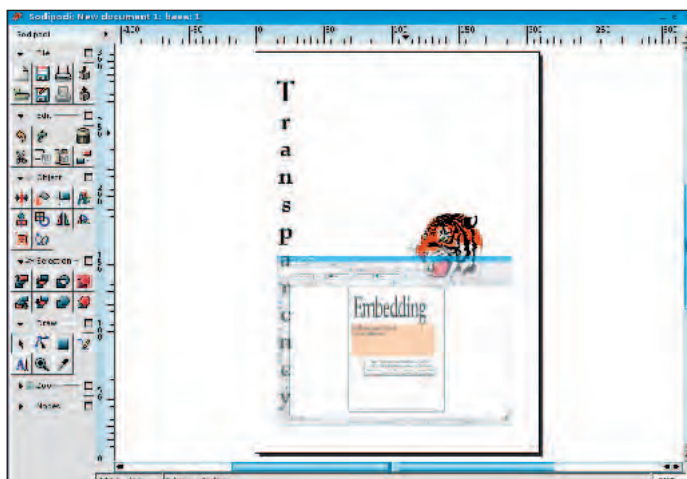
■ **VERSION** 0.34 ■ **WEB** www.sodipodi.org/

Sodipodi is an extremely promising illustration app that unfortunately suffers in comparison to its competitors due to the way in which its quirky user interface has the propensity to clutter up the desktop with windows. Some may think that this criticism may be a little unfair, as *Sodipodi*'s UI takes its cues from *The GIMP* before the refinements that were unveiled in *The GIMP* version 2. It may well have fitted in with past usage patterns of the dedicated Linux artworker, but even the famously tortuous *Blender* has managed to clarify its interface recently, so why not *Sodipodi* as well?

First impressions are that in order to get anything done, we must first get past the UI, but after a couple of hours at work (especially in tandem with *The GIMP*) it actually becomes quite usable, and right-clicking access to

almost all of the tools and options at least makes up for the wandering toolbox, which tends to get lost, especially if you're working in more than one application. We did eventually find an option to 'dock' the toolbar on the active window and this made things a lot easier, but we still suffered the occasional forlorn attempt at opening files via a non-existent menu at the top of the window.

Everything from text handling, alignment and distribution, and object transformation tools are exactly the same as *Inkscape*. *Sodipodi* only lacks the boolean operators we mentioned earlier and a couple of output options. It's worth mentioning the fantastic transparency options available in the application. It is possible not only to apply transparency effects to the building blocks of our image, but also to imported bitmap images. This



Looking for DTP ability? You'll have to forgive some of the UI quirks!

makes creating complex layered images and even DTP-style layouts quite simple.

The decision on whether you opt for *Inkscape* or *Sodipodi* depends largely on UI preference. The former is simpler in use, while the latter may feel more comfortable for hardcore users of *The GIMP*. It is easy to become accustomed to this way of working, but we still found *Inkscape* to be a more productive package.

LINUX FORMAT VERDICT

FEATURES	6/10
EASE OF USE	6/10
DOCUMENTATION	0/10
PERFORMANCE	8/10

Get over the user interface and you have the makings of quite a decent all-round illustration application.

RATING **7/10**

Sketch/Skencil

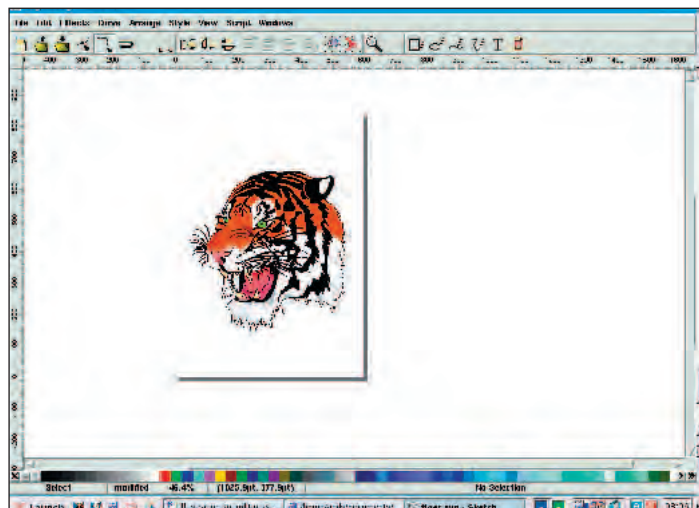
Python-built, with a new name and a retro feel.

■ **VERSION** 0.6.16 ■ **WEB** <http://sketch.sourceforge.net/>

This has a decent selection of tools if you can get it running! We had difficulties on many of the distros we tried (including SUSE's own pre-built binary), but eventually got a working

installation, via the official Debian repository, on Xandros 2.0. SVG import also needs Python-XML installed.

In visual terms, *Sketch* reminds us a little of the old paint packages on the



BBC Micro: square, grey buttons with jaggy black icons on them and menus that require a hard click to open the flyouts. This doesn't make it bad, just a little old-fashioned. The usual tools are on offer, though the creation of stars and spirals is not done 'interactively' as it is in the other packages, but using a script/dialogue approach. This makes for more precise use, but is less spontaneous.

The selected colour palette spans the bottom of the screen (reminiscent of *CorelDraw*). It takes up little space and lets the user change the colour of an element very quickly, and this simplicity in approach is mirrored in the rest of the app. Basic tools are immediately to hand, while more complexity is buried in menus; to add a gradient or pattern fill to an object involves the opening of the 'Fill' dialogue and hitting the 'Apply' button. There are no transparency options as yet, but *Sketch* does break the mold by offering a Blend tool which can 'tween' between two different shapes and colours. No other app here does this (though it was a key feature in *CorelDraw*

It's not pretty, but *Sketch* does the basics, and it does them well. More ambitious designers will probably need to go elsewhere, though.

3), which is a surprise. It's of limited use, unless you're creating lots of animated GIFs, but welcome nonetheless.

The output options include both EPS and PDF. Both performed well, as did the import filters, though we did have a few issues bringing in raw SVG data. Despite the look, *Sketch* is a good, simple vector editor which handles most tasks well. While it is stable in use, the installation difficulties we experienced probably would put most casual users off, and the text tool is very primitive, failing, for instance, to deal with line returns. DTP functions are definitely out of the question. Also anti-aliasing is a distant dream for both drawings and text which adds to the old-school feeling.

LINUX FORMAT VERDICT

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

An old-fashioned looking application with one unique tool, rubbish text handling and a few cutting-edge export options.

RATING **6/10**

ILLUSTRATION SOFTWARE

THE VERDICT

This is really a two-horse race. For general office or home use – designing presentations, fleshing out homework with clipart or the odd text effect – *OpenOffice.org Draw* is impossible to beat. It combines an extensive featureset, an intuitive interface and integration with the rest of the suite that makes cross-application work a joy. For more specialised tasks such as DTP or technical illustration, *Draw* also fits the bill. It can't be long before the other applications gain the ability to output Macromedia's Flash formats, but at the moment *OpenOffice.org* is unique in these features. The occasional annoying bug and the odd omission of an SVG import filter aren't devastating and are more than made up for by the polish of the rest of the application in most cases.

For more design-intensive tasks, *Inkscape* is the surprising winner. Although this is a young project, there is obviously ambition to push the application forward and adopt the features that may tempt the professional designer from his or her blinkered loyalty. How far a small Open Source development team can go in

The state of SVG

Scalable Vector Graphics


Vector images are perfect for display on websites, primarily because it is far more bandwidth-effective to send a stream of coordinates (this circle is 600 pixels wide, red, with an alpha value of 55) which can be interpreted by the web client/browser, rather than raw picture data. The difficulty here is that the success of the format rests on its adoption by those who develop web browsers, PDA/phone clients and other access points. If SVG doesn't gain the reach of something like Flash or Javascript, in terms of available clients, it will probably remain a niche technology. The clock is ticking on the format, with the already established Flash on one side and Microsoft's own XML-based graphical language due to be part of the next generation of Windows.

For the inveterate tweaker, those who want ultimate control, or just the curious, both *Inkscape* and *Sodipodi* have a text-based SVG viewer/editor built in: which is a bit like the code view/design view split that is available in the likes of *Quanta* or *Dreamweaver*.

Browser developers can take two approaches to supporting SVG. *Internet Explorer* and *Mozilla* are happy to hand off support to Adobe which has created a plugin for both browsers, though this obviously relies on there being a clear

incentive for users to download and install the applet. The *Mozilla/Linux* version was recently updated (at last!) and is now on par in terms of features with its Windows cousin. The *Konqueror* team, in contrast, is aiming for an embedded solution which has the benefit of offering a seamless experience for the user and should also lead to adoption of the format for general desktop use, ie for things like icons.

GNOME is already SVG-aware, as anyone who's used the excellent Scalable Gorilla theme by Ximian can attest.



Changes that are made in the XML editor are reflected back on the image.

achieving this objective remains to be seen, but the project leaders certainly deserve kudos for their imagination. They have a lot to do, and they need the support (both moral and financial)

to keep on keeping on. The presence of both *Sodipodi* and *Inkscape* here really does show the strength of Open Source development (and also adds strength to our initial argument about

a user's level of loyalty to a certain way of working). The developers of *Inkscape* were satisfied with the tools available in *Sodipodi*, but thought the interface was getting in the way of the



TABLE OF FEATURES

Name	Version	Licence	PDF Import/Export	SVG Import/Export	Bitmap support	Transparency	Text handling	Hyperlink/Imagemap	Rating
<i>Inkscape</i>	0.38	GPL	no	yes	yes	yes	good	Hyperlink	8/10
<i>Karbon14</i>	1.3.1	GPL	no	yes	yes	yes	v. good	yes	6/10
<i>OOo Draw</i>	1.1.1	LGPL/GPL	export	export	yes	yes	excellent	yes	8/10
<i>Sodipodi</i>	0.34	GPL	no	yes	yes	yes	good	Hyperlink	7/10
<i>Sketch</i>	0.6.16	GPL	export	yes	yes	no	poor	no	6/10

ROUNDUP Illustration software

LINKS

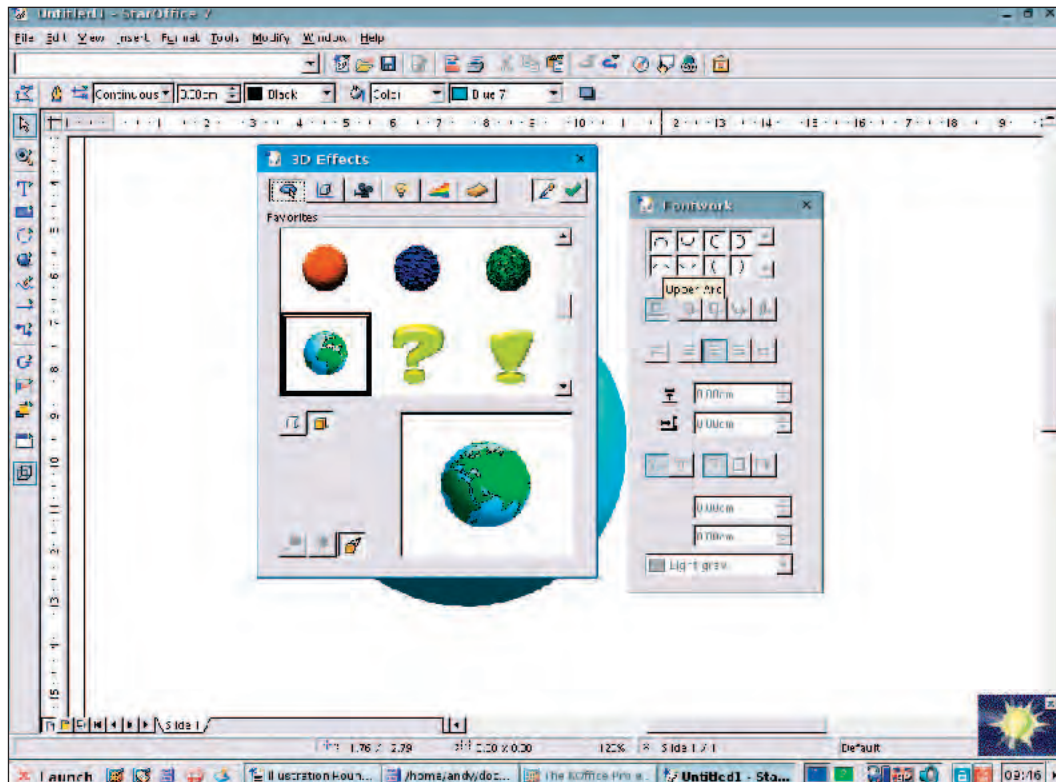
Full specification of the SVG standard:
www.w3.org/TR/SVG/

Home of the latest release of Adobe's SVG plugin:
www.adobe.com/svg/viewer/install/main.html

Home of the KSVG project which oversees the integration of the technology in KDE:
<http://svg.kde.org/>

creative use of the app. Instead of arguing and eventually coming up with a fudge, they forked the project and worked to their own goals. And while *Inkscape* has inched ahead on the features front (and in terms of interface), *Sodipodi* – and others – can, and probably will, benefit in the near future from its innovations.

Karbon 14 has the potential to be a really good illustration application. Unfortunately, it needs to get over its propensity to fall over while doing routine tasks. The developers should also concentrate on improving text handling, perhaps bringing the text tools of *KWord* on board, and also in the filters. Outputting a native *GIMP* (XCF) file with layers intact would indeed be a fantastic feature.



OpenOffice.org's illustration capabilities are the best all round, but there are unique functions in most other apps.

And finally, despite its ugly exterior, *Sketch/Skencil* is a pretty usable application. It also has one thing that others don't – shape/colour blending – so it's useful having around just for

that. As a general drawing application though, it loses out through the lack of bells and whistles. There's no transparency, for instance, and very poor text handling that seems unable

to deal with even a line return. Where it scores highly is in the area of node editing, which is superb, and its ability (though not always reliably) to output EPS and PDF files. [LXF](#)

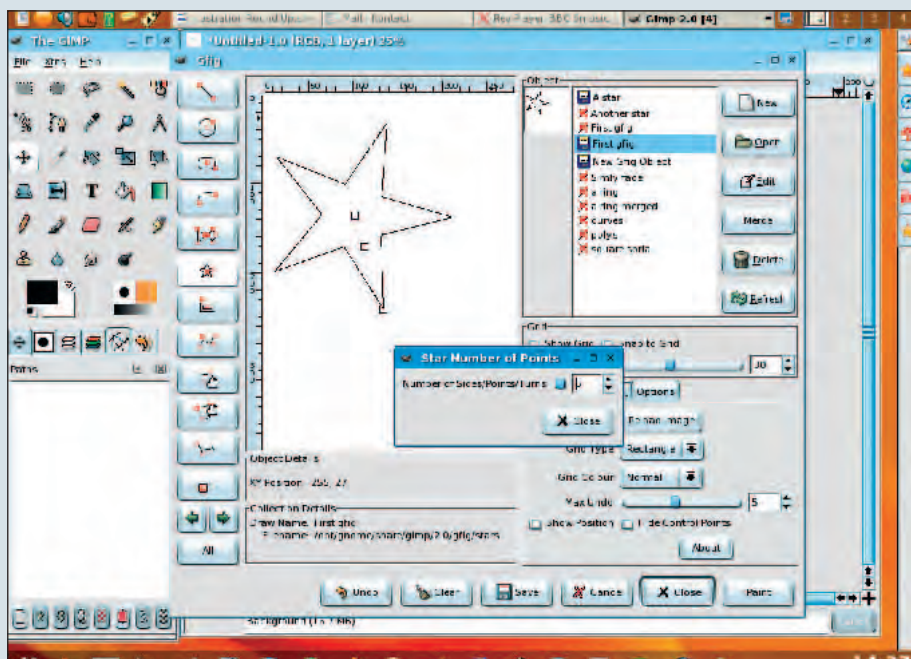
Vector art in The GIMP

Some handy hints to improve your scalable images

The GIMP, the grand-daddy of Linux art applications, shouldn't be left out of the equation when it comes to SVG art. Firstly, *The GIMP* is able to import and rasterise SVG files, though we did find that results get worse as images become more complex, and introducing text throws it almost completely. It's a better policy, we found, to allow the original application to render the file as a bitmaps through the export menu. This at least ensure the fonts and gradients will come out as intended.

Secondly, *The GIMP* can use bezier paths within the main application. This means we could build up an image composed entirely of saved paths, which would be scalable until we opted to stroke and fill them. Making an image in this fashion is not a quick process, but it does have the advantage of bringing the smooth, clean lines that vector art is known for to your bitmap creations.


Finally, *The GIMP* has its own vector applet called *Gfig*. Accessed via the 'Filters>Render' menu, *Gfig* comes complete with a selection of predefined shapes and a range of primitives. Images are built up using these primitives and then 'painted' onto the currently selected document. Again, this brings the look of vector art to a project, and is perfect for creating stars and spirals etc, but it's not an ideal solution.



Gfig offers limited vector tools in the ubiquitous GIMP.

Hot Picks

The best new Open Source software on the planet!



Mike Saunders

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

This is the place where we get to profile some of the hottest software around. Each month we trawl through the hundreds of Open Source projects which are released or updated, and select the newest, most inventive and best for your perusal. Most of the *Hot Picks* are available on our coverdiscs, but we've provided web links if you want to make sure you have the very latest version.

If you have any suggestions for software that you think we should cover, email us at linuxformat@futurenet.co.uk with "Hot Picks Request" as the subject-line, or contact us by post through the address on the *Mailserver* pages.

HOT PICKS AT A GLANCE

Aytm	48
Elmo	49
ASSP	49
Drivel	50
Vnc2swf	50
POPsearch	51
Usermin	51
Ghextris	52
openMSX	52

HOT PICKS AWARD

Everything covered in our Hot Picks section is

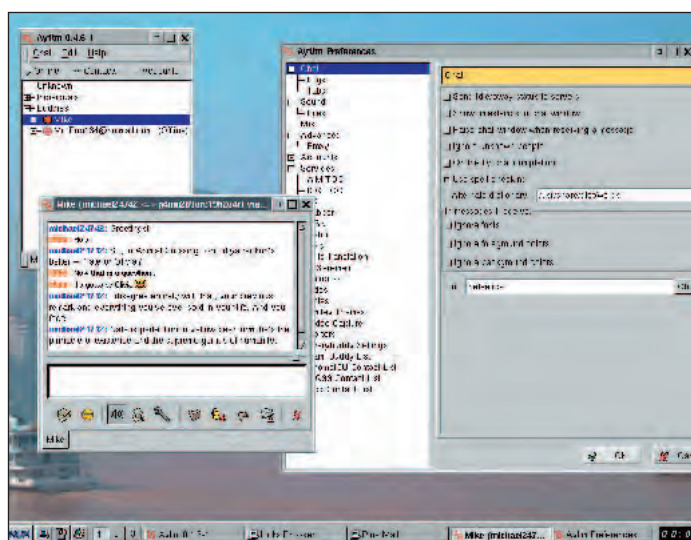
unmissable, but every month we'll be singling out one project for outstanding brilliance. Only the very best will be chosen!



INSTANT MESSAGING CLIENT

Aytm

■ **VERSION** 0.4.6 ■ **WEB** <http://aytm.sourceforge.net/>



Aytm in action – buddy list, chat window and the preferences box.

Instant messaging is incredibly popular, with millions of the world over enjoying real-time communication with friends, family and work colleagues. Many major networks and portals – MSN, AOL, Yahoo! among others – provide their own messaging systems, and some are now offering native Linux versions of their chat clients as well. We looked at a bunch of the best IM tools in LXF36's *Roundup*; although GAIM came out on top, others weren't far behind. Can newcomer Aytm dethrone GAIM as the king of IM apps?

Strictly speaking, Aytm isn't new, as it's based on the *Everybuddy* code. That app has split into two distinct projects: *EB-lite*, an interface-agnostic rewrite, and the Aytm client on test here. The name itself stands for "Are You Talking To Me?" (no doubt Robert De Niro would approve) and the client's goal is simplicity and friendly operation.

Aytm is built around the GTK1 toolkit; while this means it lacks some of the cosmetic goodies of GTK2-based clients (such as AA fonts), the upside is

that it's significantly lighter and faster in the UI department. Compilation from source is straightforward, with the ability to enable extra protocols via the `.configure` script, and we've provided binary packages on our coverdisc for various distros.

When first started, Aytm pops up the main buddy list window alongside an accounts box in which you can enter usernames and set their auto-login status. It's a clean and approachable design, and insulates the user from more advanced settings – those are accessible via the Prefs box. Aytm's buddy list is an equally frill-free affair. Different views can be enabled via radio buttons, and contacts are listed in a tree view.

It's good to talk

For chatting, Aytm offers an innovative tabbed layout to contain multiple conversations in a single window (a great saver of screen real-estate), with crisp GNOME-style tooltipped buttons assisting. Power-

users may find the quick access to certain functions an advantage over GAIM, which hides more of them in its menus, and the themeable smileys are a nice touch too. Sound events are catered for; the default set is soft and pleasant, while others can be added and specific events disabled.

Network support is one of the most crucial features in modern IM clients, and Aytm does a sterling job on this front. Like GAIM (and Trillian on Windows), Aytm works with a number of IM protocols – so you don't have to run separate clients for each network. The app handles MSN, AIM, TOC, Yahoo!, ICQ and Jabber, covering all the major players, and also includes IRC support for group chatting.

In addition to (limited) file transfer, logging and per-contact note-taking capabilities, Aytm also boasts *GnomeMeeting* and Yahoo! Webcam support (notably missing in GAIM's main branch at present), a spelling checker and an on-the-fly language conversion facility. This can be a handy addition for Open Source developers working with non-English speaking contributors, although – as with most computerised translation systems – it's less than perfect, and we had trouble getting it to work very effectively.

A handful of superficial goodies have been thrown in too, including a rainbow-coloured text filter and '133t' speakifier. More useful are the import options, which allow you to take your contact list from other clients (GAIM, GnomeICU, Licq and ICQ99), but this needs some polish. Documentation is impressively thorough, and for regular chatters the wealth of options in the Prefs box should cover just about everything – even if it can be a bit daunting for newcomers.

Aytm is a highly capable IM client with a strong featureset; it's also lighter than GAIM (generally using around half the RAM) and suitable for older boxes. And while the focus on making it newcomer-friendly is admirable, GAIM is also striving for a similar goal – so Aytm, with its abundance of prefs and features, could establish a solid position as the power-user's favourite.

CONSOLE EMAIL CLIENT

Elmo

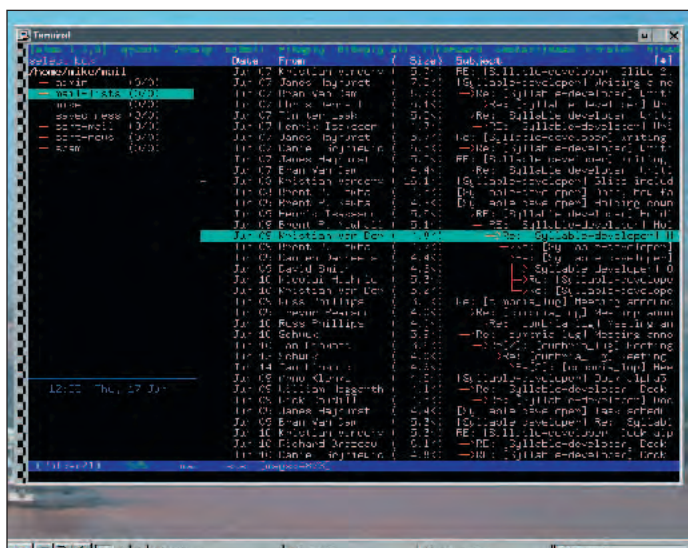
■ VERSION 1.2.0 ■ WEB <http://elmo.sourceforge.net/>

Despite the abundance of graphical email clients for Linux, with some top-notch showings from the likes of *Mozilla Thunderbird*, *KMail* and *Evolution*, text-based applications such as *Pine* and *Mutt* still retain a healthy following. They're fast, chock full of features and work brilliantly over slow *Telnet/SSH* connections. *Elmo* tries to bridge the gap, by providing a comfortable and easy-going UI whilst still maintaining text-mode speed and low resource requirements.

Kudos to the developers for making a healthy range of binaries available. Packages for most of the major Linux distributions have been put together (we've included them on our coverdiscs this month), but building from source shouldn't require any hoop-jumping, as it requires nothing out of the ordinary.

At first, *Elmo* resembles an intriguing combination of graphical and text-based apps. The keystroke guide along the top is similar to *Mutt*, and a status line sits at the bottom. Remarkably, *Elmo* supports basic themes: there's a less-cluttered variant for 80x25 ANSI terminals, along with an *Outlook*-esque layout more suitable for large xterms and framebuffer or SVGa mode consoles.

Elmo's main goal is to deliver all the required functionality in one app. So, as opposed to some other text-based clients that need an MTA to send mail, or *Fetchmail* to receive mail, *Elmo* includes its own SMTP and POP3 facilities. Everything is specified in the single `~/.elmo.rc` file, which can be configured interactively via a script (but we'd like to see a *Pine*-style in-built setup screen). *Elmo* can work with both *maildir* and *mbox* formats, and includes



Colourful threaded message antics aplenty in *Elmo*'s list view.

MIME and partial GPG support. It even features a basic Bayesian spam filter to combat the eternal UBE menace. Also, a helpful little tutorial is supplied which zips through the keybindings and general usage.

It's not perfect; opening large *mbox* mailboxes is very slow, and occasionally the program exhibits

bizarre behaviour (lockups and massive CPU usage). However, the attention paid to providing a familiar interface within the constraints of text mode deserves much respect, and the in-built SMTP and POP support gives it a boost over the competition. Great for speed and use on older machines – it's definitely one to watch.

SMTP PROXY SERVER

ASSP

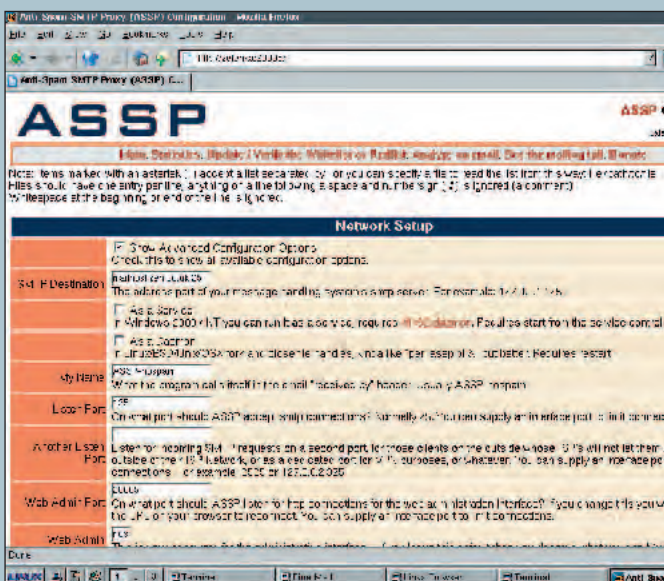
■ VERSION 1.0.10d ■ WEB <http://assp.sourceforge.net/>

Efforts are finally being made to curtail the enormous spam problem, – perhaps half-heartedly in some people's opinion – but at least it's being recognised as a major headache by governments and large companies. Some email programs do a superb job on the client-side for cleaning out junk mail, but they still require the user to download time-wasting mail in the first place. ASSP intends to be a drop-in spam-killing solution on the server-side, and stands for "Anti-Spam SMTP Proxy".

Written in Perl, ASSP thankfully doesn't rely on all sorts of esoteric modules and should work straight away on all Linux distros. Little tweaking is required before starting it up; initial steps are detailed in the rather massive (459K) HTML documentation file, along with stacks of configuration pointers and FAQs.

Pleasingly, ASSP sports a web-based configuration front-end with distinct sections and appropriately documented options – the text config file is more terse, but still hackable if need be. There's an Advanced Options toggle for revealing the more complex workings of ASSP, and it's all polished and approachable, so any moderately experienced admin should have no trouble tuning and maintaining it.

ASSP implements a number of features for dealing with spam: automatic whitelisting keeps track of outgoing email (so anyone a user contacts won't be blocked), Bayesian filtering 'learns' to identify common spam words as time progresses, and regexps can be set to mark mails and remove attachments on-the-fly. And although a robust and powerful anti-spam system can take a while to build up as the Bayes database and



Configuration is straightforward with ASSP's web-based front-end.

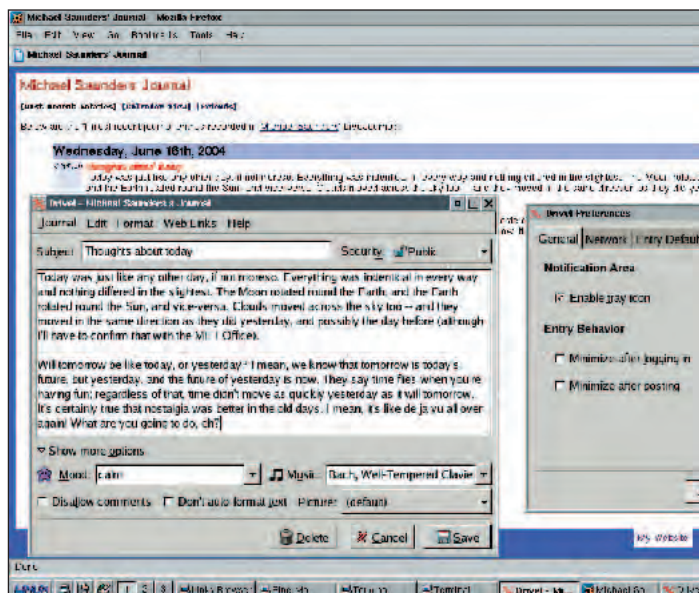
whitelist grows, in our testing ASSP performed well early on with decent success and no false-positives.

As with all spam-busting software, ASSP's effectiveness depends on the volume and type of mail passing through the system; its versatility should help it cope with most

common workloads (the coders reckon up to 100,000 messages per day is doable). ASSP has received wide acclaim and has an active support community, so if you need a quick and easy drop-in tool for cutting down spam traffic at your site, it's well worth a try.

LIVEJOURNAL BLOG CLIENT

Drivel

■ VERSION 1.0.0 ■ WEB www.dropline.net/drivel/

Drivel's entry box and preferences dialog, with the active page in Firefox.

Few people spend regular time on the Net could manage to escape the gigantic 'blogging' (web-logging) phenomenon. Used by everyone from horse trainers to kernel gurus, blogging systems are an ideal way to jot down thoughts and keep the world informed of your musings and feelings. LiveJournal (www.livejournal.com/) is just one among many successful blog sites providing a blogging service, and now, thanks to Drivel, you can blog away from a slick GNOME app without having to fire up a browser.

Eagle-eyed readers may recognise the project's URL; yes, dropline.net is best known for its fantastically well-constructed Slackware GNOME bundle. This application is built around GNOME, and requires an up-to-date distro because of its *GTK 2.4.x* dependency – failing that, you can upgrade *GTK* by hand. The RPM should work on most new installations.

Drivel initially presents a tidy login screen, spruced up with a cat staring into some trees (it must be infectious – see LXF40's *Hot Picks* for more feline fun in *Junkie*). Naturally, password-saving and auto-login are catered for.

Once recognised by LiveJournal, the main window appears and you can begin typing a new entry straight away. Like many GNOME apps as of late, Drivel puts simplicity and friendliness first but also includes a toggle for more details.

The menus provide the usual editing facilities and HTML insertion for formatting, together with the ability to save drafts, browse previous entries, add polls and manage friends. Mood, music and pictures can be assigned to the current entry, and its security level (who can view it) is tweakable too. Excellent work – all common blogging requirements have been coded in. Drivel's Prefs box is GNOMEishly minimal; you can set some defaults and choose different blog servers (eg DeadJournal, a darker workalike), but that's about it.

Drivel is polished, stable and gets out of the way – the ideal app. Just about everything a regular blogger would need is included, and it's much quicker and more elegant to use than the usual web interface. If you're a regular on LiveJournal, give Drivel (or an equivalent such as LogJam or KLuJe) a try.

SCREEN RECORDER

Vnc2swf


■ VERSION 0.4.2 ■ WEB www.unixuser.org/~euske/vnc2swf/

Screen recording apps have been doing the rounds for many years, primarily for demonstrating and teaching software and operating systems. These essentially capture the on-screen action as it progresses – writing it out to a video file or similar – and have been successful in creating tutorials, guides and even games solutions. Vnc2swf brings this capability to Linux, making use of the VNC protocol and Macromedia's Shockwave Flash file format.

To build Vnc2swf you'll need version 0.2a of the Ming SWF generating library (later versions don't work): it's provided on our coverdiscs. From there, the usual **./configure, make** and **make install** (as root) routine will put it into place. You'll also need the VNC remote-desktop software; this is bundled with most distros, or can be found at www.realvnc.com.

Before starting a recording, a VNC session needs to be active – so providing you can view it with *vncviewer* or similar tools, you should be able to capture the desktop right away. Entering **vnc2swf output.swf hostname:x** (where x is the display number) brings up a window, and tapping the **F9** key starts the recording process. The resulting SWF file is written to disk when Vnc2swf is terminated.

We were impressed with Vnc2swf as it's a doddle to use, and the output files can be viewed immediately in a standalone player or web page with the file embedded. In terms of size, the produced SWF can be tweaked with a number of parameters – frame rate, buffering and clipping – and these can help to reduce it to modern-friendly dimensions. Handily, MP3 files can be added to the generated SWF; this



Detailed backdrops and rapid window movements are a no-no if your priority is compact files – a difficult aim for graphically intensive games!

is useful for tutorials where spoken instructions are necessary.

Vnc2swf does exactly what it says on the tin: it's trim, fast and easy to set up, and the resulting animations can be viewed by over 90 per cent of Internet users. We'd love to see Vnc2swf in

wider use for teaching and advocating Linux – as not only do newcomers feel more comfortable watching an OS in action, but regular users with questions can also learn more from studying real-life solutions. A smart little app with plenty of possibilities for use.

FILESYSTEM SEARCHING POPsearch

■ VERSION 0.0.9 ■ WEB www.popsearch.net/

Filesystem search tools are generally limited to filenames and extensions. That's fine if you know exactly what you're looking for, but when you're hunting down some particular info on a subject, it gets more complicated. *POPsearch* attempts to provide "information synchronisation for the technical community" by creating a personal search engine that peeks inside files.

Don't be put off by initial glitches – *POPsearch* can be a little troublesome to set up. Alongside some dependencies that are standard in most distros, a number of additional components are required (*libwm* and *libxslt* being the most prominent). Similarly, the large C++ files push the compiler's limits – so if you get out-of-memory errors during compilation, tweak the Makefiles for less aggressive

optimisations. Hopefully this will be cleaned up in future releases.


With *POPsearch* installed, the first essential step is to build a database of files to be examined. *popbuild* runs in the background indexing files, and you can specify certain directories to spider. From there, running the main *popsearch* binary fires up a web server on port 2001; this is the tool's front-end for browsing and searching files.

Basic operation is straightforward enough – the smooth design and layout works well, although it doesn't run at a great pace. You can browse files and emails, enter and save search terms (with assistance from a thesaurus or spelling checker, which is very handy) and add comments and ratings. Popular files work their way to the top, while textual content is displayed in excellent little thumbnail-like frames.

Depending on configuration, *POPsearch* can examine the contents of *MS Word*, PDF, Rich Text and *LaTeX* files, among others. One particularly slick feature is the Research section; here the user can organise together files, mails, websites and searches into distinct topics, ready for later reference.

At this early stage, *POPsearch* is still somewhat rough around the

edges and plenty of tidying up needs to be done. We bumped into the odd bug and glitch here and there, but on the whole the system ran evenly. *POPsearch* does well in delivering a fresh approach to managing information, and has great potential for improving everyday work – we'll be keeping a close eye on this project.

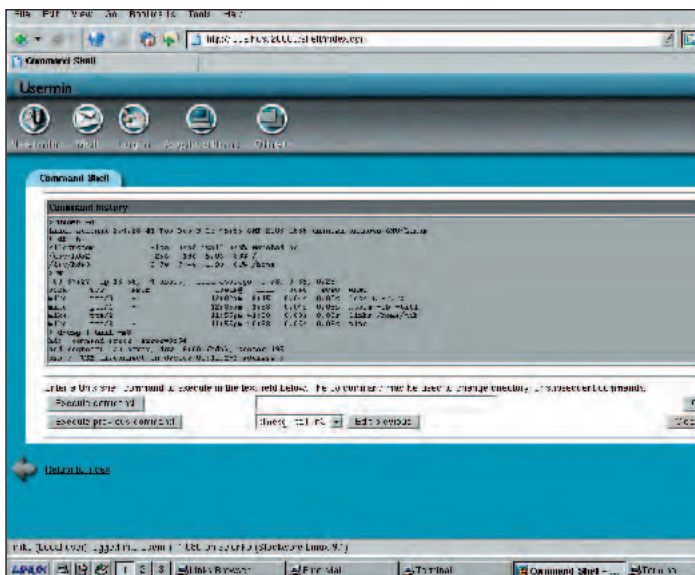


The front page of *POPsearch* looks like that of many popular search engines, but note the tiny thumbnail-like view of text files.

ADMINISTRATION TOOL

Usermin

■ VERSION 1.080 ■ WEB www.webmin.com/index6.html



Along with the Java terminal, *Usermin* provides a single-command mini-shell, making things simpler than in *Webmin*.

Long gone are the days when Linux users had to wrestle with config files and command-lines in order to get a system working properly. Thanks to the efforts of distro vendors, we now have a top-notch range of tools: Mandrake's *Control Center*, SUSE's *YaST* and Red Hat's system-config-* utilities. However, these are distro-specific, and therein lies a problem; information from one Linux installation doesn't always apply to the next. *Usermin* aims to provide a hassle-free admin system for desktops and small servers.

Usermin is from the same stable as *Webmin* (see LXF36, p30) and as a result it's similar to set up and operate. Installation from the tarball couldn't be any easier – just extract, run the *setup.sh* script and answer a few questions when prompted. While some config files are placed in */etc* for later upgrades, *Usermin* itself runs from its own directory, which is great to see. There are too many tools that claim to make your life easier, yet go through a tiresome setup procedure...

Once it has been started, *Usermin* fires up its simple web server, allowing the user to access it through a browser

(<http://localhost:20000>, or the port specified). A number of themes are supplied to tweak the app's overall appearance. Various aspects of configuration and administration are made available via the top panel; these include a quaint little mail reader, *Fetchmail*, *Procmail* and *SpamAssassin* config sections, user details and more.

The user can upload files, browse man pages, set at/cron jobs, mount filesystems and perform all manner of tasks. Some components need a full *Webmin* installation to work correctly, but most are very well done indeed (considering the browser limitations). A column-sortable process manager is available, as is a nifty Java terminal emulator for *Telnet* and *SSH* sessions.

Usermin is ideal for Linux newcomers who want to learn more about the OS without being tied to distro-specific configuration tools, and it's also highly useful for budding administrators maintaining small servers. Dedicated users and gurus will prefer the versatility of *Webmin*, but its little brother supplies a comfortable set of system tools through the familiar face of a browser.

MSX EMULATOR

openMSX

■ VERSION 0.4.0 ■ WEB <http://openmsx.sourceforge.net>**Attempts over the years to**

standardise a console format have faired badly, with the 3DO being the most striking example. Allowing a bunch of manufacturers to extend and elaborate on a 'standard' usually leads to chaos – incompatible features are added, competition becomes overly aggressive, and (as in the case of the 3DO) there's no single company pushing the machine. The first MSX appeared in 1983 as the start of a family of Z80-based home computers, growing reasonably popular in some parts of Europe and Asia, and it suffered for similar reasons. Some might argue that its close association with Tony James' outrageous pink-haired 1980s art-rock-advertising scammers Sique Sique Sputnik probably did the MSX few favours either...

openMSX is an MSX emulator which 'aims for perfection,' according

to its coders. The MSX scene is thriving due to the games released rather than serious software; still, *openMSX* will run both ROMs and


disk-based programs. The compilation process is effortless and well-documented, and there are no unusual dependencies required – just SDL and Tcl 8.4.

Boasting a number of unique features and system support functions over similar MSX emulators, *openMSX* supports a large array of virtual hardware and runs at a good pace. Compatibility is fairly strong; we encountered problems with some

games, but when it works it runs like a charm. The emulator will load .cas tape images, cartridge ROMs, .xsa disk images and others.

Along with the joystick support, full-screen mode and screenshot features, *openMSX* sports a speed-up key for skipping intros and a miniature console. Regular shell users will appreciate its tab-completion, and an impressive range of commands and options are available for tweaking the emulator (many of which are chiefly for developers and testers). The documentation is among the best we've seen for any emulator – it's concise, relaxed and very well-written, and explains the more complex aspects with care.

openMSX is one of those (increasingly common) Open Source projects with highly passionate developers. It's clear that they *love* the machine, and the attention to detail in terms of code, design and docs deserves applause. There's still some way to go before spot-on compatibility is achieved, but for now, anyone craving the classics of yore could do far worse than put *Love Missile F111* on the stereo and check out *openMSX*.



Top-down MSX racing game *F1 Spirit* running in the emulator, with the snazzy translucent console open.

PUZZLE GAME

Ghextris

■ VERSION 0.9.0 ■ WEB <http://mjriki.fi/software/ghextris>


If it ain't broke, don't fix it. That applies just as much to video games as anywhere else, although from some disastrous sequels we've seen over the years, it's clear that companies often try to milk too much from a near-perfect idea. Take *Tetris*, for example. Ostensibly simple, outrageously addictive and yet hair-tearingly difficult later on, *Tetris* appealed to just about everyone; it didn't stop aggravatingly convoluted follow-ups such as H2O's *Terisphre* appearing, though.

Ghextris, then, is a variant on the classic block-based formula with hexagons in place of squares. It's written in Python with a GNOME (GTK2) front-end, so you'll need a few supporting packages (*python-gtk*, *python-gnome* and *python-glade*). The developer has created Debian

packages and we've supplied a generic tarball on our coverdisc – it should be a breeze to install.

Cosmetically, *Ghextris* is about as rivetting as watching over-ripe fruit decompose while John Virgo laconically commentates. At five in the morning, in November drizzle, at a bus stop in Slough. Lush presentation has never been an issue with *Tetris*, though; detailed backgrounds and hyperactive music distract from the Zen-like concentration required. So *Ghextris* delivers a standard grey background on which the shapes are manipulated. Phew.

In *Tetris*, your goal is to arrange the falling shapes so that complete lines are constructed, at which point the lines disappear. *Ghextris*' gameplay is similar, but the six-sided shapes add



Bah. I started off poorly, got worse later on, and ended up in this mess!

complexity and mandate a new approach. You're constantly having to think far in advance; the blocks don't line up neatly in vanilla *Tetris* fashion, so it can all seem a bit nightmarishly challenging early on.


Sadly, there's no title music, no sound, no high scores, and not even the familiar flash feedback when lines disappear. That's not a big deal for a

desktop puzzler – but some sugar coating is always pleasant to see. *Ghextris* isn't as immediate as plain old *Tetris*, and die-hard fans won't find much in this new style of tessellating tomfoolery, but the game deserves some credit for being playable in its own right – something the three billion *Tetris* variants have rarely achieved. An engaging lunch-break toy. **LXF**

CHOOSING A LINUX FILESYSTEM

faster filesystems

Which filesystem is right for you? **Richard Drummond** has the answers.



Choice. That's the oft-touted advantage of Open Source software in general, and Linux in particular. With Linux, you have plenty of choice over which desktop environment you use, which browser you surf with, and the word processor with which you compose your text. But choice is not something that is just restricted to applications: you may also choose the filesystem that manages the data stored on your hard drive.

This may seem a little alien to users converting from the proprietary Windows and Mac worlds (where the operating system is generally pre-installed on some default filesystem and most users are content to make do with what they're given – even when a choice exists). But filesystems – like any other piece of software – differ in what they actually offer the user: even when they are designed to fulfil broadly similar goals, different filesystems offer different levels of performance, efficiency of storage, security and compatibility with other software and other platforms. Thus, the ability to pick the one that best supports your needs is important. After all, a filesystem that performs

NATIVE FILESYSTEMS

Your root choices for Linux

Second Extended Filesystem (ext2)

<http://e2fsprogs.sourceforge.net/ext2.html>

The *de facto* standard for much of Linux's life, this filesystem offers good performance and proven reliability. Thanks to its prevalence, ext2 is widely supported by tools and has excellent cross-platform support, with drivers or tools for accessing ext2 filesystems being available on many different types of – Unix-like and other – operating systems. It is non-journalled, so following a crash filesystem consistency must be checked with *fsck* (which can be a time-consuming operation for large disks).

Third Extended Filesystem (ext3)

<http://freshmeat.net/projects/ext3/>

The ext3 filesystem was created by Steven Tweedie to add journalling to the well-proven ext2 filesystem and thus greatly reduce system downtime following a crash. Although not the first or the fastest journalling

filesystem for Linux, ext3 is the most popular because of its forward- and backward-compatibility with ext2. ext3 simply is ext2 plus a journal file, so ext2 filesystems can be simply converted to ext3 and *vice versa*. It supports journalling of meta-data and optionally file data.



Hans Reiser

ReiserFS

www.namesys.com/

ReiserFS, named eponymously after its creator Hans Reiser, was the first journalled filesystem commonly available for Linux (it has been included in the kernel since 2.4.1) and, unlike ext3, it was designed from scratch to be journalled. ReiserFS is generally considered to be faster than ext3, and although shipped as default with distros such as SUSE and Linspire, it is arguably less popular due to a lack of backwards compatibility with ext2, fewer support tools and less cross-

platform support. An interesting feature of ReiserFS is the ability to pack multiple file tails (the last block in a file) into a single block to use disk space more efficiently.

XFS

<http://oss.sgi.com/projects/xfs/>

XFS was created by SGI for its IRIX Unix flavour as a high-performance and scalable replacement for the earlier EFS. The Linux port was released as Open Source in early 2000. Despite being only officially merged with the 2.4 series kernel with v2.4.25, it has been shipped as an option with many popular distros for quite some time. XFS makes extensive use of balanced trees, has native support for extended attributes and boasts unique features such as supporting guaranteed bandwidth access to a filesystem.

Journalled Filesystem (JFS)

<http://oss.software.ibm.com/jfs/>

JFS was originally created by IBM in

the mid-1990s for its AIX Unix flavour running on its RS/6000 servers, but was also available for OS/2. IBM ported JFS to Linux and released it as Open Source in 2000. It only officially became part of the Linux kernel with release 2.4.20, and the Linux port is less well-proven than its AIX counterpart. JFS supports only meta-data journalling and currently lacks support for filesystem quotas on Linux.

Reiser4

www.namesys.com/

Still in testing, Reiser4 is a radical filesystem design based on Hans Reiser's research into filesystems. It boasts many advanced features such as extensibility via plugins, fully atomic filesystems transactions and efficient storage of small files, and it employs a 'wandering logs' system to avoid the usual performance penalty incurred by the journalling of file data. It's not part of the official kernel tree yet, and has also not yet been deemed suitable for production use.

well for desktop use may not be the most suitable one on which to host a mail server, for instance.

Many Linux users are not taking advantage of their ability to choose, however. Linux distributions also typically default to using a particular filesystem and – similar to users of proprietary systems – Linux users may blithely accept that default without considering their options.

In this article, we will compare the features and performance of the current crop of Linux filesystems with the aim of educating users about what choices they offer.

First, however, we'll explore some basic filesystem concepts.

What is a filesystem?

Permanent storage devices such as hard disks are remarkably dumb when viewed at a high level. They simply support the reading and writing of fixed-sized blocks of data (blocks are 512bytes on most current hard drives). The management of such storage space is the job of operating system components called filesystems.

A filesystem provides the means to store, retrieve and organise data on

some storage medium, typically a hard disk. A disk-based filesystem sits between the applications producing and consuming data; and the device driver that performs the raw reads and writes of blocks to the device. Note, however, that applications don't usually talk directly to a filesystem. For example, in Linux, the kernel presents an interface for filesystem access, which applications normally use via the standard C library. Linux and most other Unix kernels support multiple filesystems, and rather than providing a separate kernel interface for each one, a common API (Application Programming Interface) is provided by a software layer called the Virtual Filesystem Switch (or VFS). This software layer decides which filesystem module should handle a particular access request.

'Filesystem' is also the term applied to the set of data managed by a particular filesystem software component. With a hard drive, it is normal to divide up the disk into logical slices – called *partitions* – where each partition is a subset of the disk's total number of blocks. Typically, each partition will contain a separate filesystem, but with technologies such

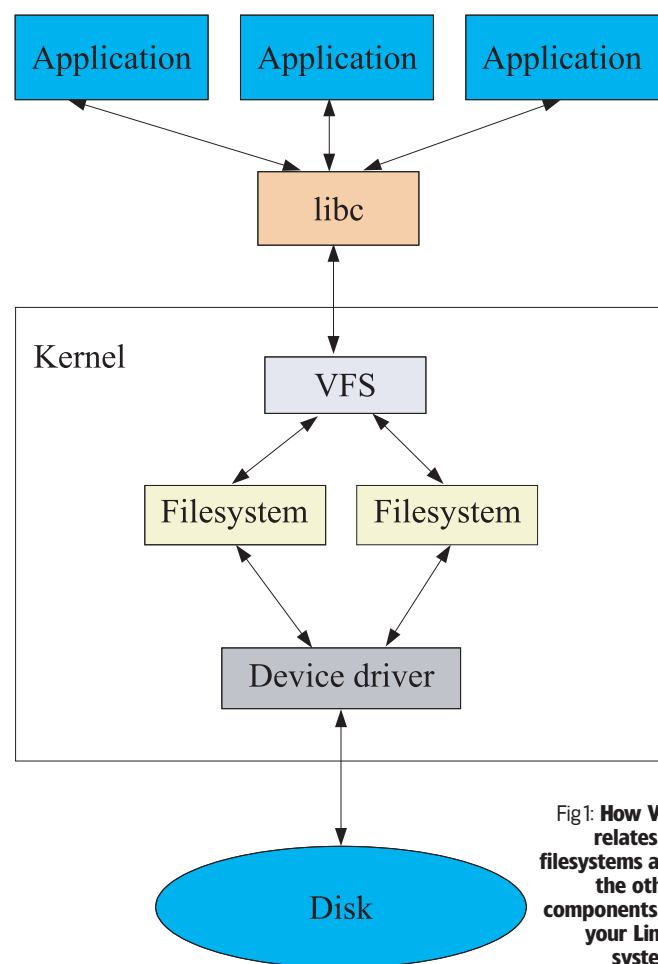


Fig1: How VFS relates to filesystems and the other components of your Linux system.

CHOOSING A LINUX FILESYSTEM




Fig2 A typical filesystem 'tree' structure, showing the path /home/rich/myfile.



as LVM a filesystem may span multiple partitions (for flexibility) or, via RAID, can span multiple disks (for increased speed or security of data).

Files and directories

The unit of data in a filesystem is the file. A file is stream of bytes, stored on one or more disk blocks, and typically meant to be accessed sequentially from beginning to end (although some filesystems may allow some structure to be imposed on file data).

To allow you to locate a particular file, each file has a filename that is associated with it (a file usually will have other properties too, such as the date it was last modified and an owner). Also, in order to help manage complexity and to allow you to more easily find a file, filesystems implement containers which are variously called directories, folders or drawers. Whatever the term, the principle is the same: a directory can contain files and other directories.

Thus, a filesystem looks to the user like a tree-like structure with files as the leaves of the tree (see Fig2 above). The beginning of the tree (or more correctly, the starting point for navigation) is called the *root*. A file's 'path' is the sequence of directories (starting at the root for an absolute path) that must be traversed to locate the directory that the file is stored in followed by the file's name itself. Each file must have a unique path. Thus in Fig2, we have the a file path /home/rich/myfile, which locates the file 'myfile' located in the directory 'rich', itself located in 'home'. Note that the use of the backslash symbol to denote the root directory and to separate the

elements constituting a file path is a purely device of Unix-like operating systems (other systems may denote his differently).

All current, general-purpose filesystems employ this basic hierarchical model of folders and files, but may well vary in all other respects. It is easy to see why: filesystems abstract away the complexities of how data is stored on disk, so each filesystem is free to cook up its own scheme for doing so (as long as compatibility with some other filesystem is not a goal). Thus, filesystems differ in the on-disk data structures and algorithms that they employ for allocating disk space and for locating and retrieving files.

From the user's point of view, filesystems will trivially differ in the restrictions they impose on such details, such as the maximum length of a filename, length of a file, numbers of files and total size of a filesystem (or such limits may be fundamental to the operating system itself); but they may also differ in more basic and discriminatory ways.

One such aspect is the range of additional properties (or attributes) that a filesystem can remember for each file. Such attributes are often called meta-data (that is, data beyond what is stored in the file itself). All filesystems will associate a name, a size and a modification date with each file, but depending on the demands of the operating system, each file may

also have access rights, a file type, and the name or ID of the application the created it. Multi-user systems will have some concept of ownership of files and also employ some means of access control. Other systems will implement 'extended attributes', allowing any software to associate arbitrary application-specific meta-data with each file (see *Extended Attributes* box on opposite page).

Unix-like operating systems in particular make some requirements of their filesystem in order to support multi-user access and to affect the 'everything is a file' philosophy of Unix. We'll examine Unix filesystem semantics next.

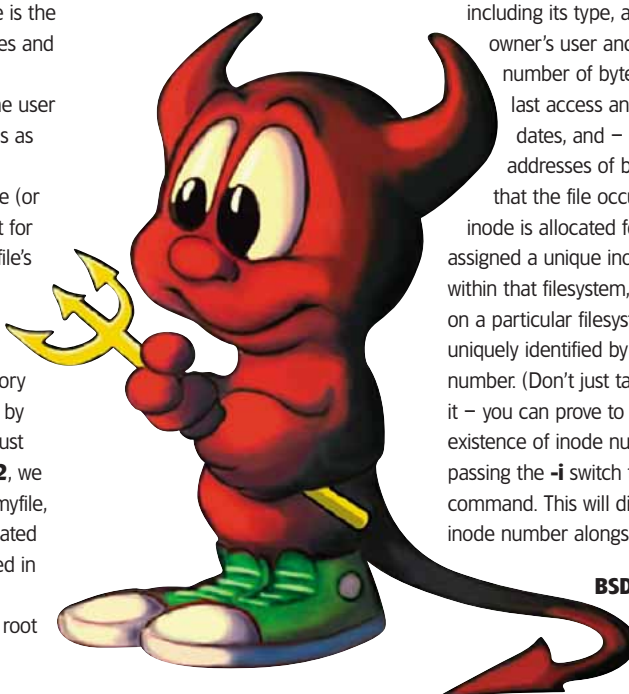
Filesystems and Unix

All current Unix filesystems can trace their heritage – if only in spirit – to the original Unix File System (UFS) and especially its high-performance successor: the Unix Fast File System, created at Berkeley in the mid 1980s for the BSD distribution of Unix. Linux's ext2, for instance, was clearly modelled on FFS, but other Unix filesystems – even those that depart more radically from the FFS design, such as the new Reiser4 – employ the same basic concepts. One such example is the inode.

In a Unix filesystem, each file is represented by an on-disk data structure called an *inode*. An inode stores the file's metadata, the basic properties of the file it represents – including its type, access mode, the owner's user and group IDs, the number of bytes in the file, the last access and modification dates, and – of course – the addresses of blocks on the disk that the file occupies. When an inode is allocated for a file, it is assigned a unique inode number within that filesystem, and so each file on a particular filesystem may be uniquely identified by its inode number. (Don't just take our word for it – you can prove to yourself the existence of inode numbers by passing the **-i** switch to the **ls** shell command. This will display each file's inode number alongside its filename).

BSD was developed at University of California, Berkeley.

BSD Daemon Copyright 1988 by Marshall Kirk McKusick. All Rights Reserved.



Remember we said that in Unix, *everything* is a file? Well, in Unix, a directory is just a special file whose contents associate the members of the directory with inode numbers. Interestingly, Unix does not specify *how* a directory should be implemented. Some filesystems use a simple linear list of associations, but a more advanced solution is to use some kind of search tree, since this offers vastly better performance for directories with large numbers of entries. (ReiserFS, XFS and JFS all employ variations of the balanced tree, whereas recent ext3 versions optionally support red-black trees for directory indexing.)

Multiple directory entries can point to the same inode. This allows the hard link, a mechanism that allows a single file to be accessed via multiple paths. An inode keeps a count of the number of directory entries that reference it. Creating a link increases that count; deleting a link decrements it. The inode itself is only deleted when its link count is zero.

Unix also supports symbolic links, a special file containing the path of the file being linked to. Symbolic links can reference directories, files on other filesystems or even non-existent files; hard links cannot.

The other special type of file in Unix is the *device node*. A device node is represented by an inode only; it has no file data. It is simply the means via which a piece of software makes I/O requests of a kernel device driver. Look at your `/dev` directory!

One other feature of Unix related to filesystems is that Unix – and hence Linux – links all filesystems in use into a single directory hierarchy. The act of opening a filesystem for use is known as ‘mounting’, and to mount a filesystem you must associate the filesystem stored on a particular device to the place in the directory tree where it will be attached, its mount point. The filesystem which is attached to `/`, the root of the tree, is known as the root filesystem. When any additional filesystem is mounted, its hierarchy of directories becomes a sub-tree of the root filesystem at the mount point.

Linux filesystems

When Linux Torvalds began developing the Linux kernel, he used Minix as scaffolding to get the project

up and running. Naturally, then, Linux started life with support for only the Minix filesystem. However, Minix is a Unix clone designed to teach operating system theory – it is not a production operating system. Thus, the Minix filesystem has some serious limitations – such as a maximum file size and filesystem of 64MB. Linux was always meant to be a practical system though, and such restrictions clearly would not do.

The first attempt to alleviate the limits was the Extended Filesystem, created in April 1992. Written from scratch for Linux, it improved on the Minix filesystem by allowing filenames of up to 255 characters and filesystems up to 2GB. It offered weak performance and lacked some features required of a Unix filesystem and so was replaced by the Second Extended Filesystem in early 1993.

Another early contender to replace the Minix filesystem was the Xia filesystem. This was based on the reliable Minix code, but alleviated some of its limits, allowing 248 character filenames and a filesystem size of up to 2GB. However, despite being more robust than the upstart, buggy ext2 filesystem, Xia faded into obscurity due to lack of maintenance (and was later removed entirely from the kernel).

ext2, of course, matured and became the *de facto* standard Linux filesystem, and deservedly so. What its detractors criticise that ext2 lacks in scalability and availability, it more than makes up for in pure speed in general applications, especially when it is compared to traditional Unix filesystems.

Today ext2 is no longer the only choice for a high-performance native Linux filesystem. ReiserFS was the first competitor to appear, initially as patches to 2.2 series kernels and then as a fully-fledged part of the 2.4 series kernels. ReiserFS was written from scratch for Linux to address the shortcomings of ext2, particularly scalability and – thanks to its support for journalling – recovery times following a crash (more information on this to follow).

ext3 was the next to appear, adding journalling to the already-proven ext2 filesystem and offering full compatibility with ext2. Later, ports of two enterprise-class Unix filesystems

EXTENDED ATTRIBUTES

Adding application-defined properties to files

Filesystems have traditionally supported only a fixed set of attributes per file: name, size, modification date, and so on. This is incredibly limiting for some types of application, so a recent trend in filesystem design has been to allow application-defined properties – so-called ‘extended attributes’ – to be stored with each file’s system metadata. The benefit is immediate and obvious. Instead of each application having to invent its own way of associating its own metadata with a file, the filesystem can do it.

Consider a jukebox app that stores track and artist data as file attributes; consider a mail client that can store the read or unread status of each mail using file attributes; or a file manager that can use attributes to store a file’s type, icon imagery and the name of a default tool to open it with. Such properties belong with a file rather than in some ad-hoc extension to the file, or worse – locked away in a separate file. Extended attributes remain with the file when it is

moved or copied, and they can be queried with a set of standard tools, rather than many third-party tools.

Linux support for extended attributes is immature. All filesystems in the 2.6 series kernel support extended attributes, but for 2.4 kernels this functionality must be applied via patches. Moreover, Linux applications are yet to support extended attributes in any meaningful way. In fact, current filemanagers are likely to *lose* extended attributes when copying files.

One important exception is Access Control Lists (ACLs). This is means for effecting fine-grained access policies for Linux files rather than being restricted to the clumsy Unix ‘owner’, ‘group’ and other permissions. ACLs in Linux are implemented using extended attributes and require kernel support and the appropriate user-space tools to manipulate those ACLs.

For more information about extended attribute and ACLs see <http://acl.bestbits.at/>.

“The unit of data in a filesystem is the file. A file is a stream of bytes that is stored on one or more disk blocks.”

were to be made to Linux: SGI’s XFS from IRIX and IBM’s JFS from AIX. And if that’s not already enough choice, the successor to ReiserFS – Reiser4 – will soon be stable enough for real use, and brings cutting-edge filesystem features to Linux.



CHOOSING A LINUX FILESYSTEM

« Crash recovery

The most critical requirement of any filesystem is reliability. Speed is important too, but a fast filesystem that loses your data is clearly of no use. While the ext2 filesystem has proved to be remarkably robust over the years, it does have one important flaw: it makes no guarantees as to the

“The most critical requirement of any filesystem is reliability: a fast filesystem that loses your data is clearly of no use.”

state of the filesystem should your computer crash while writing to the filesystem. If this happens with ext2, the filesystem is left in an inconsistent state and must be verified before it can be used again.

How, then, does Linux know that a filesystem is inconsistent? Well, the filesystem implements a status flag.

When a filesystem is mounted with write access, it is marked as unclean; if unmounted cleanly or remounted as read-only, then it is marked as clean again. If you try to mount an unclean filesystem as writable (perhaps because the system crashed without cleanly dismounting it), Linux will warn you and mount it as read-only. It is then necessary to run the filesystem checker, *fsck* (the Unix equivalent of Microsoft's dreaded *scandisk*), to check the filesystem for inconsistencies and repair any damage. (The ext2 filesystem also employs a mount count which for safety forces a filesystem check after it has been mounted a certain number of times. This forced checking period can be set with the *tune2fs* tool.)

Each filesystem has its own checker, which will be, by convention, named *fsck* with the filesystem type as a suffix. The *fsck.ext2* tool is very good at its job and will rarely fail to recover a filesystem. The problem is that since *fsck* has no idea what area of the filesystem was being accessed when the system went down, it has no option but to scan the entire filesystem for consistency. This means the time taken to check a system is proportional to the size of the disk. A large filesystem may take tens of minutes or even hours to validate – and this down-time may be completely unacceptable for a server, and is certainly frustrating for desktop users. What's more, this problem is worsening all the time, since technology is advancing disk size more rapidly than the speed of disk access.

Dear Journal...

The most common solution for removing the lengthy filesystem consistency check following a crash is for the filesystem to keep a record of recent filesystem activity, and thus always know how far it is from a consistent state. Such a record is called a journal or log and gives rise to the term a 'journalling' filesystem. Scanning a journal and returning the filesystem to a consistent state is known as 'replaying' the journal, and, since the journal is vastly smaller than the size of the filesystem itself, replaying is a quick operation.

The concepts behind journalling are borrowed from the world of

databases where consistency is paramount. A journalling filesystem tries to ensure that all modifications to the filesystem are done atomically. Such an indivisible modification is called a transaction. A transaction usually corresponds to higher level filesystem operations such as creating, deleting or renaming a file.

Two types of filesystem journalling are possible: filesystems that perform journalling on meta-data and those that journal both meta-data and file. We'll consider the former for now, since it's easier to illustrate.

Any filesystem transaction will require that a file's meta-data be updated. To effect an atomic update, a meta-data journalling filesystem first records the updated meta-data in its journal and then marks the transaction as complete in the log. It then updates the file's actual meta-data, typically by copying it directly from the journal. When that operation has finished the filesystem marks the transaction as complete in the log and removes it from the journal. If the system goes down at any time during the transaction, recovery is simple. The crucial thing to realise is that the only possible inconsistency occurs when copying the updated meta-data blocks from the journal out to the filesystem for real – this is the only time that the on-disk data structures are modified. If a crash occurs here, consistency is restored when replaying the journal simply by copying the blocks out to the filesystem again. If a crash occurs at any other point in the transaction, no action is needed and the journal entry can be safely ignored when replaying.

Note that a filesystem which journals meta-data makes no guarantees that file data won't be lost – only that the disk can always be quickly returned to a consistent state. It is possible to journal file-data as well, but since this means that all file data has to be written twice: once to the log and once to the filesystem, so it obviously causes a significant hit to performance. Meta-data journalling, in contrast, is cheap to perform and in some cases may even speed up filesystem writes since it allows meta-data changes to be batched together. Currently, ext3, ReiserFS and Reiser4 all support data journalling, but only

FOREIGN FILESYSTEMS

Accessing filesystems from other platforms.

AFFS

The Amiga Fast File System (no relation to the Unix Fast Filesystem) has been used for both floppy and hard disks on Amiga computers, and has seen support in Linux for some time (notably the Amiga was the first non-x86 platform Linux was ported to). The Linux AFFS module has had sketchy maintenance over the years, however, and lacks support for features in newer variations of FFS such as international filenames and long filename support.

BeFS

The Be Filesystem was introduced with Release 2 of BeOS. It is an advanced filesystem, with support for large file sizes, extended attributes and, uniquely, live indexing of file attributes. Linux support for BeFS is still experimental.

HFS, HFS+

The Hierarchical File System is the native filesystem for Apple's older Macintosh computers. Extended HFS (or HFS+) was introduced with Mac OS 8.0 and is still used in OS X, but has only recently gained support in Linux (a fact that caused pain to Mac Linux users). HFS+ is an advanced filesystem, supporting both the bizarre separation of files into resource and data forks

required by the classic Mac OS and Unix semantics, and more recently journalling. Linux support is still largely experimental.

MSDOS, VFAT

The Linux MSDOS filesystem implements FAT12 used in earlier MS-DOS version and still used on floppies, and the newer FAT16 and FAT32 variations, which were implemented to overcome filesystem size limitations. The VFAT filesystem adds support for long filenames (introduced with Windows 95), where the MSDOS modules employs the older 8.3 filename scheme. Besides compatibility with Microsoft, the VFAT filesystem is widely used for removable media such as Flash cards, so support in Linux is essential.

NTFS

The NT Filesystem was created for Windows NT, Microsoft's attempt to create an enterprise class operating system free of the legacy of DOS, and it is now the default filesystem of Windows XP. Linux support for NTFS has traditionally been poor, and the kernel driver in fact is still semi-functional. More advanced Linux support is available by way of patches to the standard kernel (see <http://linux-ntfs.sourceforge.net/>).

Reiser4 claims to do so without a major impact on performance – thanks to its ‘wandering logs’ system. For this reason, both ext3 and ReiserFS both default to only journaling meta-data.

The agony of choice

We have now discussed the behaviour and features of filesystems in some detail. So how do you go about picking the right filesystem for the job? Clearly, that’s not a very easy question to answer.

Two significant criteria are speed, and the existence of support of tools for a particular filesystem; both issues we will discuss more fully later on in this article. For just now, we’ll look at more general issues.

First of all, it’s worth underlining that all the Linux filesystems we’re examining here implement broadly similar features. They are all suitable for use as your root filesystem and will perform adequately as such, especially if your requirements are fairly modest. Thus, if you already have an installation that performs satisfactorily with its current filesystem, then you should think hard before switching to a new filesystem. With the exception of ext2 and ext3, converting a system from using one filesystem to another is not a trivial undertaking. You can’t convert directly from one filesystem to another – you’ll have to copy files and directories from the old filesystem to the new. The archiving tool *tar* is your friend here; but still, this is not a task to be undertaken lightly.

Another point that should be considered carefully is whether your computer system will be able to support the new filesystem without having to make any major changes. If you’re installing a new Linux system, then this is not an issue. Current versions of major distros have wide filesystem support with Debian, Fedora Core, Mandrake and SUSE all featuring support for the filesystems we’re discussing (apart from the new Reiser4). Upgrading an existing installation to a new filesystem, however, may require rebuilding your kernel with support, plus adding the necessary support tools.



With those caveats out of the way, let’s address the issue that we’ve all been waiting for: speed.

Speed tests

So which Linux filesystem really is the fastest? In the fine tradition of computing, it depends. To try and gauge the relative performance of the current Linux filesystems, we performed a series of tests on each one using the *Bonnie++* benchmark. *Bonnie++* implements two classes of filesystem test: I/O tests which measure throughput and CPU load when accessing a single large file, and

file creation tests that measure how fast a filesystem can create, access and delete files, first in sequential order and then random order and the associated CPU load each time. (see www.coker.com.au/bonnie++/ for more information about *Bonnie++*).

We tested using an Athlon 1200 box with a 1GB of RAM running a 2.6.7-mm4 kernel (the ‘mm’ tags denotes Andrew Morton’s kernel tree). This kernel was chosen simply



CHOOSING A LINUX FILESYSTEM



because the latest Reiser4 patches were available for this kernel. The drive being abused was an 80GB Maxtor unit accessed via a UDMA100 ATA adaptor.

For each filesystem on test, a fresh filesystem was created (using the same drive partition each time) with a block size of 4096KB and the filesystem's defaults for other settings, and then a series of *Bonnie++* tests performed. Throughput to a 2GB file was measured and then the file creation tests performed in turn with Obyte files, 32byte files and 64KB files. Each test was performed ten times and an average set of results obtained.

We don't have the space to print all the results from *Bonnie++*, but *Table 1* below shows an interesting selection of the results we obtained which should at least make some trends clear. The speed results only are also graphed on the previous page.

And the winner is?

The first thing to notice with the results is how close all the filesystems are in the I/O tests. There's barely a 20 per cent difference in speed between all the contenders in these tests. Not surprisingly, ext2 generally performs the best here (since it doesn't have the overhead of maintaining a journal), but ext3 and

JFS are not far behind. Interestingly, the newcomer, Reiser4, performs relatively poorly here, except in the rewrite test – where it steams ahead of the competition.

In the file creation tests, the filesystems being examined exhibit markedly different performances. Look at the tests operating on empty files. ReiserFS and Reiser4 are in a league of their own here with Reiser4 creating empty files a staggering twenty times faster than ext2. Also, with small files, both XFS and JFS put ext2 and ext3 to shame. With larger files, the difference between the tested filesystems is less noticeable.

You should be aware that it's always dangerous to draw conclusions from such a seemingly limited set of benchmarks, but you should bear in mind that we did actually perform many more tests than we have the space to print results from here. We mention this primarily because in our tests, ReiserFS and Reiser4 didn't always perform better than their competitors. One example of this was when supplying the **-b** switch to *Bonnie++*. As this excerpt from the *Bonnie++* documentation states:

"A note on blocking writes"

I have recently added a -b option to cause a fsync() after every write (and a fsync() of the directory after file create

or delete). This is what you probably want to do if testing performance of mail or database servers as they like to sync everything. The default is to allow write-back caching in the OS which is what you want if testing performance for copying files, compiling, etc."

Our experience was that if you do run *Bonnie++* with these 'blocking writes', then both ReiserFS and Reiser4 take a serious hammering in performance, and quite often produce results significantly slower than the others. For the record, XFS and JFS seem to suffer much less with blocking writes than the others.

The moral of this story is that if performance is a critical concern governing your filesystem choice, then benchmark the applications you want to run on that filesystem. While the Reiser flavours seem to out-perform the others in many tests (especially when operating on large numbers of small files), depending on how your applications are written, this might not always be the case.

Support tools

Speed should by no means be the only issue that dictates your choice of filesystem. Support tools for such tasks as doing backups, recovery and general manipulation may also be a determining factor. While each filesystem comes with its own suite of basic user-space tools, real-world tools may be sadly lacking.

ext2 – and by inference ext3 – enjoy the widest tools support. *Parted*, the disk partition manipulation tool; *partimage*, the filesystem back-up tool, and many other recovery tools all support ext2 (and generally ext3). This is not surprising, really.

What may be more surprising is the lack of support for the competing filesystems. Even ReiserFS, which has been around for donkey's years, lacks the support from third-party tools that ext2 and ext3 enjoy.

Parted This is the Open Source answer to commercial partition management tools such as *Partition Magic*. It has only fairly recently gained support for creating, copying or non-destructively resizing Reiser partitions (and then via a non-standard library). Support for XFS and JFS (and of course Reiser4) is non-existent.

TABLE 1: BONNIE++ BENCHMARKS

I/O Test	Units	ext2	ext3	ReiserFS	Reiser4	XFS	JFS
Write sequential blocks	KB/sec	41841/20	41109/23	40158/24	35017/18	39585/17	41820/19
Rewrite sequential blocks	KB/sec	17121/11	17277/11	17442/11	19185/16	18044/11	17056/9
Read sequential blocks	KB/sec	34138/13	33373/12	29832/12	33764/15	32240/11	33987/11
Random seek	ops/sec	189.2/0	176/0	142.5/0	170.4/15	177.3/0	182.8/0
Sequential create Obyte files	files/sec	1561/99	1269/99	20467/99	29223/99	3276/35	11478/31
Random create Obyte files	files/sec	1609/99	1273/98	19553/99	11283/99	3425/39	7504/50
Random delete Obyte files	files/sec	3589/99	3090/98	14502/98	10848/99	1181/15	5390/23
Sequential create 32byte files	files/sec	1503/99	1074/99	8187/99	17283/99	2952/49	7679/38
Random create 32byte files	files/sec	1521/99	1089/98	7944/99	7246/99	3121/48	4324/40
Random delete 32bytes files	files/sec	3138/99	2535/97	10832/99	8901/98	1191/17	1959/2
Sequential create 64KB files	files/sec	1088/96	686/86	-	1294/61	842/25	865/15
Random create 64KB files	files/sec	1090/95	690/86	-	656/39	851/26	684/15
Random read 64KB files	KB/sec	9912/98	7151/84	-	7082/99	10108/99	10508/99
Random delete 64KB files	files/sec	3116/97	1307/58	-	5868/99	655/15	1590/8

Notes

1. The result of each test is quoted as a pair of numbers, the first being speed in the specified units and the second the overall percentage of CPU time consumed by the test.
2. The file creation tests on 64KB files failed with ReiserFS causing a segfault from *Bonnie++*, presumably due to bugs in this particular kernel version (this test worked with ReiserFS from the standard 2.6.6 tree, for example, but obviously it would be invalidate comparisons to change part of the test-setup).

partimage How does this popular tool for making back-up images of filesystems fare? Recent versions can handle ext2, ext3, ReiserFS and JFS – but no XFS. (For more information, see www.partimage.org).

dump Maybe *dump*, a standard filesystem backup tool beloved of system administrators works across the board? ext2 and ext3 have a *dump* command. XFS has its own version of *dump*. Anybody else? Sadly, no.

gpart What about *gpart*, a tool for recovering partitions tables on x86 machines? ext2 and ext3 support, certainly, but detection of ReiserFS partitions has been decidedly iffy; and while it now purports to handle XFS, support for JFS non-existent.

As you can clearly see from our examination so far, choosing a filesystem is more than just picking the one with features or performance you need. If you're not convinced yet, one more compatibility issue may convince you.

Booting and root

If you wish to be use a particular filesystem as your root filesystem, you obviously need to be able to boot your system from that filesystem. Obviously, your kernel must have support for your chosen filesystem when it tries to mount root – either your the kernel should be compiled with a built-in handler for that filesystem, or your system boots from a ramdisk and the necessary filesystem kernel module is provided there. Just as importantly, however, your system's bootloader must be able to load the kernel from your chosen root filesystem. Whether this is possible depends to a large degree on your bootloader.

Some bootloaders, such as *LILO* for x86 machines, locate the kernel to load from disk via its block address on the disk. They don't need to know anything about the filesystem on which the kernel lives. Other bootloaders, such as *GRUB* on x86 systems and *yaboot* for recent Apple Macs, locate the kernel via its filesystem path (for example `/boot/vmlinuz-2.6.7` or whatever) and thus need explicit support if they are to work with a particular filesystem. Luckily, the latest versions of *GRUB*

have support for all the current filesystems and patches are available if you need it to boot from Reiser4. Similarly, *yaboot* currently supports them all apart from JFS and Reiser4. Users of other bootloaders (especially on non-x86 systems) might not be quite so fortunate. For example, the *quik* bootloader for older PPC machines can only boot from an ext2 partition (and won't even work from an ext3 partition that has been unmounted cleanly).

If you face such restrictions as these with your system's bootloader, a simple solution is to keep the boot directory on a separate partition, formatted with a filesystem that *is* supported. Just create a small ext2 partition (100MB should more than suffice) that is mounted on `/boot` and is thus accessible via your bootloader. If you are at all worried about corruption, you can even mount that filesystem as read-only by default, only remounting as writable when you need to install a new kernel.

Filesystems – picking and choosing your solution

At the start of this article we asked the question “*What filesystem is right for you?*” As with just about any software engineering question, this one has no simple answer.

It is possible to make some broad generalisations, however. The safe choice for any Linux installation is to use ext3 (there must be a reason that

Red Hat uses ext3 exclusively in its Enterprise distros, after all). While its performance won't necessarily impress you that much, it's an easy upgrade from ext2 and enjoys widespread support. One point worth making, however, is that in some scenarios ext2 may actually be a *better* choice. The primary advantage of ext3 over ext2 is improved recovery times following a crash. If you don't expect your system to crash, the better performance of ext2 may be more important than fast recovery.

For desktop use though, it would be hard not to recommend ReiserFS. Despite the wide ranging support, good performance and journalling are clearly a win. Xandros and Lindows use ReiserFS for good reason.

For server loads, the question becomes more interesting. Our testing would leave us to believe that XFS and JFS perform better than ReiserFS or ext3 for mail servers and databases. In fact, JFS impresses all round due to the low CPU usage we saw and its excellent throughput.

The fact is that all these filesystems have their pros and cons, and no single filesystem fits all sizes. Investigate your options, experiment, and choose the one best suits your needs based on your experience and requirements. We at *Linux Format* are eagerly waiting for Reiser4 to be declared stable: it implements technologies that advance Linux to the leading-edge of filesystem designs. At the very least, it will be a spur for the further development of all Linux filesystems. **LXF**

TABLE 2: FILESYSTEMS COMPARED

Features	ext2	ext3	ReiserFS	Reiser4	XFS	JFS
Journalling (meta-data)	no	yes	yes	yes	yes	yes
Journalling (file data)	no	yes	yes	yes	no	No
Supported block sizes (1)	1K, 2K, 4K	1K, 2K, 4K	1K, 2K, 4K	1K, 2K, 4K	512B-64KB	4KB
Directory search	linear	linear/RB-tree	B-tree	B-tree	B-tree	B+-tree
Extended attributes	yes	yes	yes	yes	yes	yes
Quota Support	yes	yes	yes	no	yes	No
Tool support						
Parted	yes	yes	yes	no	no	no
Partimage	yes	yes	yes	no	no	yes
Dump	yes	yes	no	no	yes	no
GRUB	yes	yes	yes	yes (2)	yes	yes
yaboot	yes	yes	yes	no	yes	no

Notes

1. In practice, block size is limited by the architectures pages size, for example, 4KB on x86 systems.
2. Support via unofficial patches.

What on Earth is... EM64T?

As the world hastens towards 64-bit desktops, another processor joins the race. Biagio Lucini answers your questions on Intel's EM64T.

» **There's been a lot of noise recently about a new type of chip, called EM64T. Can you tell us more about it?**

EM64T stands for 'Extended Memory 64 Technology', and is the name for the new technology Intel is introducing for extending the capabilities of the x86 (commonly referred to as 'IA32') processor family, without breaking compatibility with 32-bit apps.

» **Ah, just an extension. What's the reason for all the hype, then?**

To be fair, this extension (or rather, this set of extensions) is revolutionary in some sense: it will bring 64-bit technology into the IA32 family of processors. Hence, contrary to what the company has stated several times in the past, Intel will enter the low-end 64-bit market in a less risky way: by providing processors that can be used both in 64-bit mode and 32-bit mode. The latter guarantees that all the applications written for IA32 processors can be run on the new

processors without modification, while the newly introduced 64-bit mode will allow EM64Ts to exploit all the benefits of that more advanced technology.

» **Why is Intel presenting this technology as an extension?**

The reason is that for the first time in its history, Intel has been forced to catch up with its most aggressive competitor: AMD. While Intel was stating over and over again that there was no room yet in the low-end server and workstation market for 64-bit technology, AMD released an affordable 64-bit CPU that has obtained an unanticipated (at the least by Intel) degree of success. And while Intel was dwelling on the necessity of basing 64-bit computing on a fresh architecture to try and push sales of its Itanium, AMD had the intuition of building its 64-bit CPU on the solid foundations of the IA32 architecture.


With AMD's x86-64 CPUs gaining a significant market-share – mostly at the expense of Intel's Pentium 4 and Xeon – an inversion of policy by Intel was needed, and needed *fast*.

The main reason why x86-64 technology has received a lot of attention is that it provides 64-bit capabilities without breaking compatibility with 32-bit CPUs, which means that running 32-bit applications on those CPUs does not require modifications or recompilation. This feature has been welcomed by companies that have invested billions in 32-bit computing: by adopting this technology they can save those investments and at the same time start to move to the 64-bit world. This is a good example of how *evolution* is better received than *revolution* in the IT world. Now, Intel has probably taken note of this trend and has acted accordingly: it wants to exploit this mixed feeling of embracing a new technology without breaking compatibility. For this reason, 64-bit manufacturers have to reassure potential consumers that the new technology is not new, but just a simple extension to the well-tested IA32 architecture.

However, despite the fact that the actual implementation keeps its promise of continuity, to a greater extent it is being regarded as a new architecture, and is starting to be known as IA32e.

» **Are there other technical differences in the structures of IA32 and IA32e, besides the number of bits that can be crunched at once?**

IA32e will come with an extended cache and a faster bus; but the most apparent difference is in the number of registers: the new processors will have sixteen general-purpose registers and sixteen XMM registers (these are for SSE and SSE2 instructions), which are respectively 64-bits and 128-bits in width. Contrast this with an IA32 processor,



which has eight 32-bit general-purpose and eight 128-bit XMM registers. Since an increased number of registers means an increased amount of data that can be stored on-die at the same time, the performance of an IA32e processor is bound to be much better than that of an IA32 processor.

» Are there fundamental differences between the Intel Itanium and EM64T families of processors?

Yes: the two are based on completely different technologies and architectures, so binary files for one will not work on the other. However, there is the possibility that the current 32-bit emulation layer of Itanium will support 64-bit applications for the EM64T in the future, but it is unlikely that an emulation layer for Itanium executables on EM64T will ever be written, since (unless you want to test applications developed for the Itanium on IA32e systems) this makes very little sense.

» Why is Intel manufacturing two incompatible 64-bit CPUs?

One answer is the different needs of the market: the Itanium chip is better suited to servers with high workload, has a larger cache, the possibility of addressing a bigger amount of memory and superior performance for



sort of witch hunt, with experts looking for clues about Yamhill in processors already present on the market (if you are technically inclined, you can have a look at www.chip-architect.com and search for Yamhill). Clackamas was the next working title that the technology was known by during development, before it finally adopted the official EM64T designation.

» Which models of the EM64T processor are currently available for consumers to buy?

No EM64T processor is on the market at this very moment. The first EM64T should be Nocona, code-name of the new dual-processor Xeon, which was planned for release in the second quarter of this year, but whose introduction is probably going to be delayed. This autumn, or at the latest at the beginning of 2005, Intel should release Potomac, a multi-processor Xeon with support for at least four CPUs, and a consumer version based on the same core as Prescott (Pentium 4).

To tell the whole truth, it is not fully correct to say that there are no EM64T processors on the market today: the EM64T technology is already present in Prescott (released last February), but has not been enabled, which essentially means that it can't be used yet. You may remember that the very same thing happening with Hyperthreading technology some time ago.

software explicitly for the Itanium. With EM64T (which in principle addresses lower-scaled situations) Intel hopes to boost its 64-bit strategy. It is not the case that the same hardware vendors who were sceptical about the Itanium have quickly stated that they will fully support EM64T.

» We have heard names like Yamhill and Clackamas relating to Intel and 64-bit computing. What are they?

They are just early working names of the EM64T technology. Yamhill was supposedly a secret project by Intel to build a processor with the same characteristics as AMD's x86-64. The world should not have known about it (it was at the time when all official Intel press announcements were repeatedly stating that the only 64-bit processor it would ever support would be the Itanium), but the news slipped through and was picked up by several Internet sites. There was also a

floating-point operations.

Moreover, it scales

better on large Symmetric

Multi Processor (SMP) systems. This

makes the Itanium much better suited for

large-scale enterprise tasks. Despite this, the revolutionary architecture on which the Itanium is based (Explicit Parallel Instruction Computing) has met with a lot of scepticism, mainly because it breaks compatibility with the past. Not many companies have been willing to invest in the Itanium, hence not many software houses have made an effort for writing

WHAT ON EARTH EM64T

« » Many big companies – including Intel – are pushing 64-bit CPUs. But why should the world want to move towards 64-bit computing?

64-bit CPUs have several advantages over 32-bit ones, but they could have also disadvantages. There are several advantages in having a 64-bit capable processor. One good reason would be faster execution of high-precision calculations.

Consider, for instance, the case of double-precision mathematics: while a 32-bit processor can handle it, it has to use some tricks, which means that operations on big numbers are slower than operations on single-precision floating-point numbers. On a 64-bit processor, double-precision numbers are native, and theoretically code that use them should gain a factor of four in speed. Moreover, always thanks to the ability of dealing with double the number of bits, the range of integer numbers that a 64-bit number is capable of representing is considerably larger. This is relevant, for example, in multimedia and graphical applications, including games. Both *MySQL* and *Oggenc* make extensive use of 64-bit integers in their code.

Then there is the well-known problem of addressing memory: with a 64-bit processor, the limit is set at around 4 Gigabytes of memory, a quantity that is soon filled by high-end applications like large databases, while with 64-bit addressing the limits are pushed back to 16 exabytes. Of course, there aren't just advantages: the capability of working with more data at once also means bigger workloads on the off-processor components. Hence, to avoid bottlenecks, 64-bit processors often have a large cache and are mounted on fast boards.

« » Wow! This new 64-bit technology sounds really exciting.

Of course, it is exciting, but it's by no means new: despite the

fact that 64-bit computing has had a lot of recent publicity, 64-bit computing has been a standard in the supercomputing world for a while. High-end 64-bit capable processors made their appearance in the early 1990s. For instance, Compaq/HP Alphas have been around for more than a decade (though until few years ago, owned by a company called Digital). To be fair, there are also 128-bit processors on the market – have you ever realised that a game console like the PlayStation2 is powered by a 128-bit CPU?

As far as operating systems go, Unix has always been one of the biggest players in the 64-bit OS arena, to the point that some dialects of Unix have been explicitly concocted for 64-bit systems. Despite the fact that its origins are strictly related to the x86 world, Linux is hardly a newcomer in the 64-bit world either: the Alpha port of the kernel was officially integrated with the development versions of 2.1.x, which were the first kernels ready for 64-bit.

Nowadays the Linux kernel has strong support for 64-bit machines, having been ported to MIPS, SPARC, Itanium and other 64-bit architectures. Strange as it might seem to you, even Microsoft explored the 64-bit workstation market by releasing a version of Windows NT for the Alpha family. It was quickly withdrawn because of a lack of demand.

« » Why has it taken so long for 64-bit technology to be available outside the supercomputing sphere?

Well, this is mostly down to market demand: what was unthinkable a few years ago has become almost a necessity nowadays. One of the most evident reasons is the need to extend the maximum amount of available RAM.

« » What are the practical alternatives to 64-bit computing?

One of the obvious alternatives is multiprocessor (especially SMP – Symmetric Multi Processor) systems: having more processors on the same board overcomes many of the limitations of a single 32-bit processor. For instance, with two processors, you can execute two processes simultaneously rather than just appearing to do so. However, the main issues here are scalability and cost: a system with N processors can increase the performance of a single processor by a factor of at most N , while by incrementing the number of processing bits the potential gain becomes exponential. If we consider the fact that multiprocessor systems are by no means cheaper than 64-bit systems, we can simply conclude that migrating to 64-bit CPUs is currently the best choice for those who need to go beyond the limitations of 32-bit processors.

« » What are the differences between the various 64-bit processors available on the market?

Technically, the architecture is the biggest difference, ie the set of instructions supported by the processor

and the way in which they perform computations. This is what makes different architectures incompatible between each other. Among the various architectures are RISC (with variants like UltraSPARC, Alpha, MIPS), PowerPC (mainly adopted by IBM and Apple), EPIC (currently used only on Itanium), and x86-64 (AMD Opteron, Athlon 64, and now EM64T). These architectures have very different characteristics, and as a consequence they are used for very different tasks.

« » Is there room for 64-bit computing on every desktop?

Some would have you believe that 64-bit computing is only useful for servers, workstations and hard-core gamers. There is some truth in that: we know very well that for the average user who just wants to check his email and browse the Internet, many of the desktop systems currently available are completely feature-bloated. However, many things already happen under the hood of an operating system or a desktop manager that are meant to improve the so-called user experience, and these often require a fast processor and a large amount of RAM. In both cases, a 64-bit processor gives indisputable advantages over a 32-bit one. It is however very likely that the way to 64-bit computing for desktop systems will pass through the wide adoption of the technology in the server and workstation market. But there are some that think differently: Apple already has a series of computers based on the PowerPC 970 processor (a 64-bit chip) that targets the desktop market.

« » What are EM64T's competitors?

With EM64T, Intel will probably be competing with itself, in the sense that the target of its new processors will be the very same market segment that is now dominated by its Xeon processor. However, as we mentioned earlier, the new technology will not be deployed side-by-side with the old one, but will replace it. In other words, the new families of Xeon and Pentium 4 processors will embody EM64T. The other obvious competitor (actually probably the very reason why this technology exists at all) is AMD's x86-64: the two architectures have similar characteristics and similar scope, hence they are likely to become the most popular choices for low- and middle-end computing solutions.

« » With x86-64 already on the market, why did Intel decide to introduce a new architecture?

The reason is simple: psychology. The world of consumer microprocessors was decisively led by Intel until only a few years ago, with AMD catching up more recently. The introduction of the first Opterons last year meant that, for the first time, AMD had a clear technological advantage over its main competitor. While the virtues of Opteron were being loudly praised by many vendors and specialised publications (including this one), Intel was



left trying to convince the world that a 64-bit processor for a low-end market was not needed, while the high-end market was better catered for by its young and revolutionary Itanium. Intel responded very quickly to a clear advantage that the Opteron had over the Itanium, namely the possibility of executing 32-bit code, with the introduction of a software compatibility layer on the Itanium.

Despite the claim that this emulation layer enables an Itanium with a given clock speed to execute 32-bit code at the same speed as a 32-bit native processor with the same clock, this strategy failed. Meanwhile, it was rumoured that Intel was secretly working on a processor very similar to the Opteron. The obvious failure of the anti-Opteron campaign forced Intel to unveil that processor, which happened early this year. The Intel processor will differ from an x86-64 processor – enough of a difference not to call it 100% x86-64. However, the conceptual differences are no more than those between an Athlon XP and a Pentium 4 (for instance, the number of available registers on an EM64T and an x86-64 is exactly the same). As a consequence, software written for the one will mainly work also on the other.

» What are the main differences between the x86-64 and the EM64T architectures?

Not many: EM64T supports Hyperthreading technology and some extensions like SSE3, while x86-64 does not. Some instructions available on the x86-64 are not present in the EM64T (an example is the *3DNow!* set). There are also more subtle differences in the implementations and in some low-level instructions. This will probably give a bit of a headache to the developers that want to make use of them and grant portability of the code at the same time. These differences will force them to write variants of the code for the two different architectures. However, those are minor flaws that will be ironed out with time; what really matters at the moment is that the two processors are based on an essentially compatible architecture.

» How could that be?

x86-64 is based on open specifications: everyone is allowed to give an independent implementation of those. By using the early documentation of x86-64 and working directly on x86-64 processors, Intel was able to essentially clone (or 'reverse-engineer') the architecture supported by the 64-bit processors of AMD. However, both Intel and AMD have their very own story, to which they are tied: Intel's 64-bit processor will grant support for Hyperthreading and SSE3, but not for *3DNow!* and *vice versa*.

» What are the most visible differences between an x86-64 and an EM64T system?

That's a good question. The processing unit is not the only component of a system. As far as 64-bit computing goes, the CPU unit on its own can't guarantee speed: you also need to provide data to the processor at a reasonable speed. AMD has fixed this with the so-called Hypertransport technology, which guarantees that data can be fetched at around 6.5Gbit/sec in each bus. An Opteron comes with three buses, which increase the total theoretical transfer speed to 19.5 Gbit/sec.

Intel is answering this need by providing support for PCI Express (another technology that the company is developing) and DDR2 memory. Other enhancements included in the EM64T platform with respect to the IA32 are better power management and high-definition audio.

» What is the relative performance between EM64T and x86-64 processors?

With EM64T not yet on the market, this is really hard to quantify; there have been suggestions freely circulating on the net that EM64T performs very badly, especially if compared with the Opteron. While these rumours were surely based on beta versions of the chips that might well have used emulation layers (which would be slower than the hardware execution itself), it is possible that EM64T will need to play catch-up with x86-64. However, we would not expect it to be terribly inferior when it is released; we will be able to tell you a bit more about that over the coming months as chips hit the streets.

» Which operating systems can I run on an EM64T box?

By and large, the same operating systems that you can run on a x86-64 box. There is no problem with Linux, since the kernel has been ready for x86-64 for while now. Among vendors that officially support this 64-bit CPU, there are Red Hat (starting from Enterprise Linux 3 Update 2), SUSE (9.1 and the upcoming SUSE Linux Enterprise Server 9) and Mandrake (10.0 for x86-64). Others will follow very shortly (check out the site of your favourite distribution to see whether they have plans to support EM64T). You may have not realised this, but a 64-bit native version of Windows has not been released yet (it is expected later this year), and the likelihood is that the first EM64T processor will hit the market before a compatible OS branded by Microsoft. This will give a considerable advantage to our favourite OS.

» Can't I run a 32-bit operating system on an EM64T box?

Of course you can, but you will lose all the benefits of having a 64-bit processor: only 32-bit applications, operands and operations will be available when a 32-bit operating system is installed.

In more detail, like x86-64 processors, EM64T can operate in three modes: legacy mode, compatibility mode and 64-bit mode.

In legacy mode, the operating system is 32-bit, the processor is seen as an ordinary IA32, *ie* it has all the operational modes of an IA32. On the minus side, this means that you can't use the extended set of registers and the possibility of addressing an increased amount of memory. The advantage is that you can still use your old applications, without any need to recompile them. In this mode the processor will probably only be marginally better than a present-day Xeon.

The compatibility mode is sort of a mixed mode: you will have a 64-bit operating system, but apps still run as if they were 32-bit. This means that while the system can generally benefit




from the 64-bit extensions, applications probably won't. The mixed mode is especially useful for the transition period, when many proprietary applications won't be ready to exploit the 64-bit advantages. If you already run your IA32 in some legacy mode, you should be careful when using this mode.

Finally, 64-bit mode let you run full 64-bit apps, and obviously for that you will need a 64-bit OS. It may take some time before enterprise closed-source applications will get there. Note however, that if you run 100 per cent GPL software on Linux, you should be already able to use this mode.

» Which applications are 64-bit ready today?

Except for the Linux operating system and related GPL software, just a handful of proprietary applications have been already ported, among which are the Nvidia drivers for XFree86. However, many applications that do not require heavy intervention at low level should be pretty easy to port, since some compilers (including GCC) have already been ported (the Intel compilers do not support EM64T at the moment, but the release planned for the third quarter of this year will). To simplify the tasks of developers, Intel will also shortly release a new version of both its *VTune Analyzer* and *Performance Libraries* (MKL and IPP) with support for EM64T.

Developers interested in porting their applications can already look at the two volumes of the Developer's guide, which can be downloaded at www.intel.com/technology/64bitextensions. 

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 tutorial 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: web servers for all!

Being a beginner doesn't mean you can't use the same Linux tools as big companies – dive into the *Abyss* today! **p68**

GIMP & Inkscape


Whether you know it as vector, scaleable or structured artwork, it can be done with Linux **p72**

Scribus DTP

"It's about building compound objects." That's what you get when you let a programmer write intros – a long-winded way of saying 'boxes and shadows' **p76**

Programming with SDL: games

With fewer pictures to the page than even PHP, this issue the fish come on by the dozen. Raise your shields and get ready for some action – at last! **p78**



72

Practical PHP

Flex and hone your web programming skills with this all-new project – our comprehensive instructions make it no harder than a Las Vegas divorce **p83**

KDevelop

Signals and slots form the core of the event-driven programming model behind Qt – Sir Jono Bacon is your trusty guide on the quest for the holy GUI **p88**

TIP OF THE MONTH!

CONVERT AFRESH

If you've installed a new distribution recently, don't get in a rut using the same old processes. Why not take a look around your system and see what extra options are available?

Take Mandrake, for example. Control Center 10 has a whole section for security into which few people ever venture. And yet in there is the excellent 'Levels and checks' dialog, where you can easily set up Mandrake to run regular security checks on your files and network subsystem that can help spot if you've been hacked. It can also help lay down basic rules about security, such as the minimum number of uppercase letters. The

most important options are arguably under 'System Options', where you can disable remote X Windows sessions, disallow remote root logins (a smart move unless you really need this regularly), or even whether root should be allowed to login directly at all (as opposed to being switched to using the **su** command).

In KDE 3.2, several great new features were slipped in. For instance, under the Control Center, click 'System', then 'KDE Performance'. Here you can set up KDE to automatically preload several instances of *Konqueror* so that when you launch a web browser it appears

instantly. Another hot favourite is the new options in *Konsole*, 'Monitor for Activity', 'Monitor for Silence', and 'Send Input to All Sessions' – they all revolutionised the way we work at LXF.

The LXF team isn't blameless: we get in a rut too! So, the moral of the story is: dig through your system to take a good look at all the hard work the developers have put into it. Changing from SUSE 9.0 to 9.1 or from Mandrake 9.2 to 10.0 isn't just about getting a new version of KDE 3.2 with prettier menu transparency: there are thousands of changes waiting to make our lives easier – if we'd just take the time to find them!

BEGINNERS' GUIDE TO LINUX APPLICATIONS

Run your own web server

While a web server may at first seem a bit of a difficult subject for a beginners' tutorial, **Andy Channelle** discovers that it doesn't have to be as complicated as you might imagine.



NOTE

Throughout this *Beginners'* tutorial, for the sake of clarity (and due to the naming scheme that the *Abyss* developers have opted for), we are referring to the place where typed commands are entered (which might usually be called the console) as a '**terminal**' and the web-based GUI for configuring the web server as the '**console**'.

Linux, as many of our readers will know, is fast becoming the standard web server operating system; figures from Netcraft (www.netcraft.com) suggest that over two-thirds of web page hosts are running *Apache*, mostly on Linux. While *Apache* is a very powerful application, it can also be daunting for new users so, in order to reduce the complexity somewhat, we have opted to set up a personal web server using an application called *Abyss*. This is a small HTTP server freely available from www.aprelum.com, and though it started out as an Open Source project, later versions have been distributed under a more restricted license. It is ideal for our purposes, as it is very simple to install and use, while still having some advanced functionality, such as the ability to use Perl and PHP scripts. Crucially, it also has a very neat graphical user interface (GUI) for day-to-day management tasks, which makes everything a lot easier, whether you happen to be a beginner or not.

But why run a web server in the first place? There are a number of reasons that spring to mind. Firstly, you can use it to test your own web pages before FTPing them to an ISP's space; and while it's possible to test HTML pages without a web server, you won't get very far when your needs go beyond static pages. You might also want to make some files available to a limited number of people – family photos for instance – over a broadband connection. Running a server allows you not only to easily restrict users by adding authentication, but also makes it possible to limit access to specific times, perfect for an small office Intranet. It also removes any space restrictions imposed by most Internet Service Providers.

The first task is to acquire and install the software, so either get it from the 'Magazine' directory of this month's *Linux Format* coverdiscs, or visit the Aprelium website and download the 118K file 'abwsx1.tgz'. This file needs to be extracted into a directory where you have read and write access. It is easiest just to extract

the directory into '/home/~username/'. Once the file is extracted, open a terminal and navigate to the web server's directory with:

```
cd abyssws
```

To start the webserver, simply type:

```
./abyssws
```

A short note will be displayed explaining that an initial configuration file has been created along default lines. To test that the server is working, open any web browser on the system and, in the location bar, type:

```
http://localhost:8000
```

This simply points the browser at the locally running server using port 8000. Traditionally, web servers run on port 80, but, citing the fact that some ISPs block access to port 80 and many distributions reserve this port for use by *Apache*, the developers of *Abyss* set the application to serve on this port. This can be changed later, if needed, to 80 so that the address entered doesn't need the 8000 prefix; browsers automatically attempt to connect on 80. The server can be accessed over a network either typing its internal IP address (the address of the network card used for the internal network):

```
http://192.168.0.1:8000
```


or use the computers hostname (if it is set up):

```
http://linuxbox1:8000
```

Accessing the server outside of the network, *ie* over the Internet, involves entering the IP address of the externally facing network card or modem, see the box about *Dynamic DNS and Firewalls* over the page for more information on this.

Site and server management

Once you are satisfied that the web server is working, the first and most simple job to do is change the default page to something more personal. To do this, launch *Nautilus* or *Konqueror* in file manager mode and navigate to the /abyssws



This is the main interface for the web server. This screenshot also includes the server status screen in the foreground.

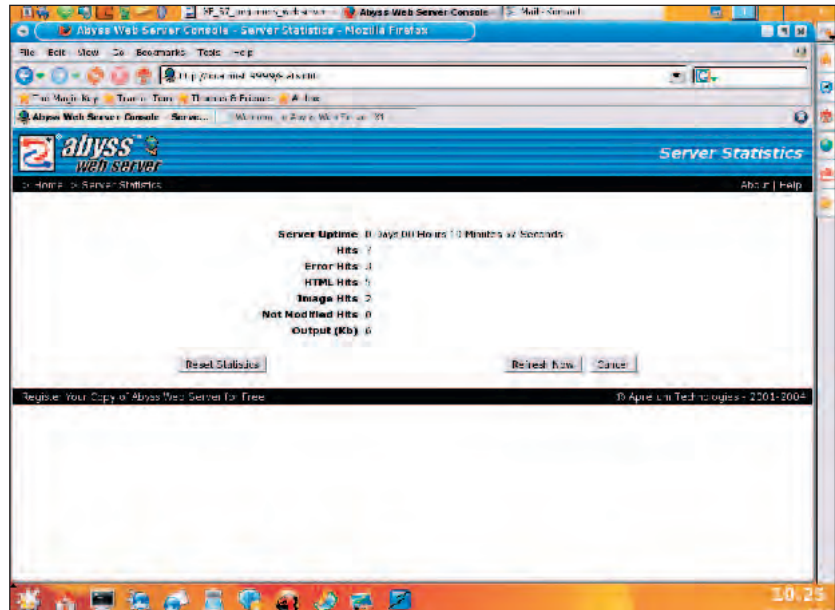
directory where the application was extracted earlier. This folder should contain six new directories and a few other files. Inside the directory labelled 'htdocs' is a couple of files, a HTML document and an image. This is the 'content' directory of the web server and anything placed in here will be accessible via the web (or internal network) when the server is running. Just like *Apache*, the *Abyss* server will search first for a file called 'index.html', so replacing this file – which is what we saw on first logging in to the server – with our own file called index.html will change the default page. All HTML files will go in this directory and as it is set up within /home it means no fussing with permissions. It is also simple to set up *Quanta*, *Dreamweaver* or some other HTML editing application to use this as the project root if you're using *Abyss* as a testing server. However, if the server is going to be exposed properly to the Internet, it is far better to work in another directory and then update pages when they are ready. While websites should always be regarded as a work-in-progress, no one wants their users to encounter half-typed sentences or rough drafts of pages.

This is a plea for order. When making a website, sorting out the directory structure beneath this 'root' folder will really pay dividends as your project progresses. At the very least create a new directory in here for the graphical elements of the site to keep them apart from the HTML files.

Once the index page has been changed and the server tested, we can begin looking at the configuration options *Abyss* provides. First, you may notice that if the terminal in which application was launched is shut down, the application itself shuts down too. To get around this, we can start *Abyss* as a daemon, which means the server runs in the background and we're free to close down the terminal. Just use:

```
./abyssws -d
```

Test the server is running as before and then access the console, the *Abyss* configuration utility, by pointing any web browser at **http://localhost:9999**. On its premiere launch, the console should ask for a user name and password (x2) which will be used subsequently to authenticate console access. Choosing a secure password is important for general computer use, but doubly so when configuring a web server. Remember that this server will (probably) be visible to any passing script kiddie with a grudge against the world, a passion for site-defacing and a hard disk full of unsavoury images, so pick something made up from a selection of upper- and lower-case letters and numbers that is longer than six digits. Don't use the word "password"!



Once past the authentication screen, you will be presented with the console, a simple icon-based page from which all server options are available. Most of the options we will be covering are under the 'Server Configuration' heading, so we'll come back to that in a moment and run through the other options.

Server Status provides options for starting, stopping and halting the server. The difference between the two latter options is that 'stopping' the server allows you to restart it from within the console, while 'halting' will mean going back to a terminal and launching from scratch.

Access Statistics tells you how long the server has been running, how many hits it has received and the amount of data, in K, served since it was started. It doesn't quite compare to the access logs available in *Apache*, but it does the job. This page automatically refreshes every 10 seconds, but you can do a manual refresh by hitting the appropriate button. There is also an option for resetting the stats.

Help and Support does as expected, providing access to the project's web forums and documentation as well as a handy link to 'Discover your IP address'.

Console Configuration allows you to change two things: the port used for accessing the console (currently 9999) and the username/password combination required to gain access. Changing the former involves restarting the server, and any change in the latter will launch the log in dialogue you encountered the first time the console was fired up.

About Abyss Web Server displays the products user and distribution licence. This is a closed-source product but is freely redistributable subject to the terms of the licence (no changes to the software and no money changing hands beyond reasonable distribution fees).

Right, back to the main configuration tasks. Hitting 'Server Configuration' launches a basic options screen. This is where the server's root folder is set, the documents path, and the port used

Windows users can only look on in admiration and envy – an uptime of over 19 minutes!


TIP

The advantage of having everything in a nice, simple directory structure is that once a day you can right-click on the root folder and copy it to a backup location. Now if your meddling with settings causes a hiccup, you can simply restore this backup – complete with working server and documents – without losing too much work.



TUTORIAL Beginners' Linux: Web Server

PHP is present and correct on the server.



◀◀ to access the server. We are going to leave these as they stand (if you want a more standard server, change the port to 80) and hit the 'Advanced' button to get into the real guts of the server.

Extended features

Instead of running through all these options available, we have set out to achieve two tasks. The first is to add support for PHP to the server so that a weblogging service can be added. The second is to create a folder that will serve content only to those with a valid user name and password.

To add PHP support, you first need to confirm PHP is present on the server's host machine. Open a terminal and type:

which php

If PHP is installed, the output will show the path to the executable, which you need to make a note of. It's usually either `/usr/bin/php` or `/usr/local/bin/php`. If it's not installed you can download and build the source yourself (from www.php.net/downloads.php) or install it using your distribution's package management system.

Once PHP support is sorted out, go back into 'Server Configuration>Advanced' and select the 'CGI Parameters' option. 'CGI Processing Enabled' should already be set to **Yes**, but if it isn't, change it.

In the table labelled 'CGI Interpreters' hit the 'Add' button. In the 'Interpreter' field, enter the path to PHP we found earlier, then add **php** to the 'Associated Extensions' field and ensure 'Automatic Update of the CGI Paths' is set to **Yes**. Hit 'OK'. If everything is going to plan we should see `/*.php` in the CGI Paths table. Click the 'Add' button in the User CGI Environment Variables and add the following entries:

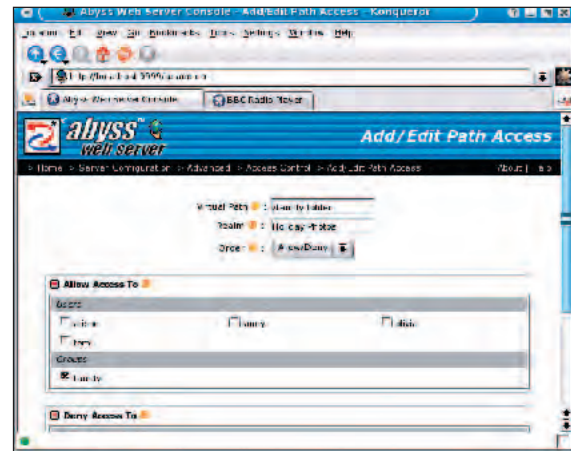
name: REDIRECT_STATUS

value: 200

Hit the 'OK' button, and then press 'OK' in the CGI Parameters



This process (with the obvious changes made) can also be used to add Perl support to the server.



All users and groups are listed in the Access Control page.


page to take you back to the main configuration screen. So far, as mentioned before, *Abyss* will recognise `index.html` (and `index.htm`) as valid index pages. However, we also want to confer this status on `index.php`. Go into the 'Server Parameters' page and click 'Add' at the bottom of the index files table. Simply type **index.php** in the text field and hit the 'OK' button. Do 'OK' right back to the top page, the one with the basic options' and then press the 'Restart' server button at the top of the screen.

Now we can add php pages, anything from a nice simple blogger app such as Bloxom (www.bloxom.com) to a comprehensive news publishing system like Fusion PHP (www.fusionphp.net/) to the `/cgi-bin` directory and the server will be able to serve them!

User authentication

To authenticate users or groups they first need to be defined. To do this, go into 'Server Parameters', hit the 'Advanced' button and then select 'Users and Groups'. Before creating Groups, it's probably best to add a few users. In the Users table, simply hit 'Add' and type in a user name and password then hit the 'OK' button. Do this for as many users you want to access authenticated parts of your site.

For example, if want to add some family photos to your site, but don't want strangers looking at them, you would first corral these users into a group called 'family'. Click on the 'Add' button in the Groups table, give your collection a name before using the radio buttons to select who belongs to this particular group. Finally hit 'OK' and then restart the server to allow the changes



Visitors will need to supply a user name and valid password to get into this site.

DYNAMIC DNS AND FIREWALLS

Overcoming the problems of shifting IP addresses

A DNS address is the set of numbers that allows other computers to connect to yours. It will consist of four numbers separated by full-stops, eg 82.32.37.167, and if you've purchased a domain name, you can redirect www.domain.co.uk to those numbers. Every computer connected to the Internet has one of these numbers (even the ones not running servers). However, most ISPs have a limited number of IP addresses to dish out, so you may find you have a dynamic IP address which changes from day to day. This makes it particularly difficult to run a web server, as your users won't know where to point their browsers.

Fortunately, there are a number of services that allow you to resolve a domain name to a dynamic IP address; they can also provide software to automatically poll your IP address and update their systems so that the domain also always ends up pointing to the right numbers.

There are a couple of services like this, but we've chosen NoIP (www.no-ip.com) because it has a no-frills free service and the service's auto-update client runs on Linux. After you've signed up for the service, which includes choosing a domain name from their lists

and adding your own subdomain (ie bob.hopto.org), you can download the software and install it on the server. It includes a binary edition (which will work with most Linux distributions) and source version which need to be compiled with **make** and **make install** (as root). It's much better to use the source version, as the script also runs through the initial configuration, a process which requires you to enter your NoIP user name and password.

When you have installed it, you can rerun the configuration at a later date by opening a terminal and typing:

```
/usr/local/bin/noip2 -C
```

Once it is all set up exactly as you wish, running the NoIP client is simply a case of typing:

```
/usr/local/bin/noip2
```

This will harvest your IP address and pass it on to the NoIP server, which will then associate it with your selected domain. Of course, you could also redirect any other domain name to resolve to your NoIP address, which in turn would connect it to your server.

Finally, you will need to ensure your Firewall is configured to accept connections on the appropriate port.



Services such as those provided by NoIP can get round the problems of dynamic DNS.

to take place. To the right of each group and user name is a pair of icons. The pencil allows you to edit the properties of an entry, while the trashcan deletes it.

Now users are added you can make the changes necessary to add restrictions to a specific directory. To start with we have created a file called 'family folder' inside /htdocs where, you'll remember, all the HTML content goes. This is where we will post our restricted documents.

Again, everything is done via the console, so do 'Server Configuration>Advanced>Access Control'. There's just one table in here called Access Control. Click the 'Add' button.

There are two elements to this page. The first job is to define the path to the restricted space and then the users or group with

(or without) access to it. To set up access to our previously created directory, the 'Virtual Path' needs to be set to '/family folder'. Setting this with a single forward slash would restrict access to the entire server. The 'Realm' provides an option give the restricted folder a description such as 'Holiday Photos' which is displayed on the authentication dialogue.

The Allow/Deny drop-down sets how Abyss manages authentication, but is not something to worry about at the moment.

Finally, there are the two Allow and Deny tables. Users or groups selected in the former table will be able to use their credentials to enter the site, those listed in the latter will not. Once these are all set, the server needs to be restarted before the changes take effect. [LXF](#)

NEXT MONTH

Firewalls have made a brief appearance this time. Next issue, we'll be taking a closer look at your first line of defence against the evils of the Internet, and a few other security matters.

TUTORIAL GIMP & Inkscape

BEZIER CURVES

Working with Beezers

Logo design is an art unto itself, and no professional graphic artist would entertain the idea of working on a logo design project without the aid of a vector illustration tool. **Michael J Hammel** shows us how Inkscape makes creating shapes a breeze, using a powerful and easy-to-use feature: Bezier curves.



Last month, we took our first look at *Inkscape*, the latest Open Source graphics tool to sweep the Linux desktop marketplace. *Inkscape* fills a long-standing void in this arena: vector illustration. Vector design allows artists to create at one resolution and easily scale – without distortion – to any other resolution. This flexibility is extremely important in many areas of graphical art, especially in logo design.

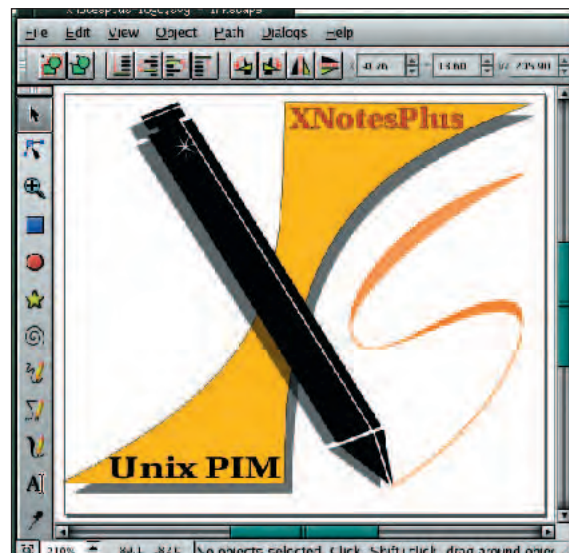
Logos are the focal point of a corporate or organisational identity. A logo can be used on letterheads, envelopes, memo pads, coffee cups, and pens in the office; it must also serve as a location marker on conference banners or building signs and as an identifying label on company vehicles.

It would be a waste of energy (and a pain to boot) to have to design multiple logos with a raster-based tool (like *The GIMP*) for printing at different resolutions. Instead, logo artists use vector illustration to design a single logo, and they then can scale it appropriately for the intended use.

In this month's tutorial, we'll recreate the *XNotesPlus* logo paying special attention to the use of Bezier curves. If you've ever used Mac OS 9 or higher, you'll probably have come across *Stickies* – an application that lets you place Post-It™-style notes on your desktop. *XNotesPlus* is a sticky notes program that takes this idea one step further into the domain of the PIM (Personal Information Manager): with over a dozen features like email, text search, printing, audio/visual alarms, date and calendar inserts, categorisation by coloured projects, etc. It also includes an interface to the 3Com Palm Pilot, allowing backups and syncing of backups, uploading *XNotes* to the Pilot and downloading Pilot

Memos into *XNotes*, and a complete interface to the *Pilot Address Book*, allowing you to use addresses from the Pilot to print envelopes and email notes to your contacts. *XNotesPlus* comes with support for both *GTK* and *Motif/LessTif*. For more information about *XNotesPlus*, see www.graphics-muse.com/xnotes/xnotes.html

Before we dive into the meat of the tutorial, we will first need a little background on terminology.



This month's finished project: the *XNotesPlus* logo in *Inkscape*.

IMPORTANT LINKS

Inkscape
www.inkscape.org/
 DAG repository
<http://dag.wieers.com/home-made/apt/>
 What are bowls and counters?
www.philfonts.com/character.html

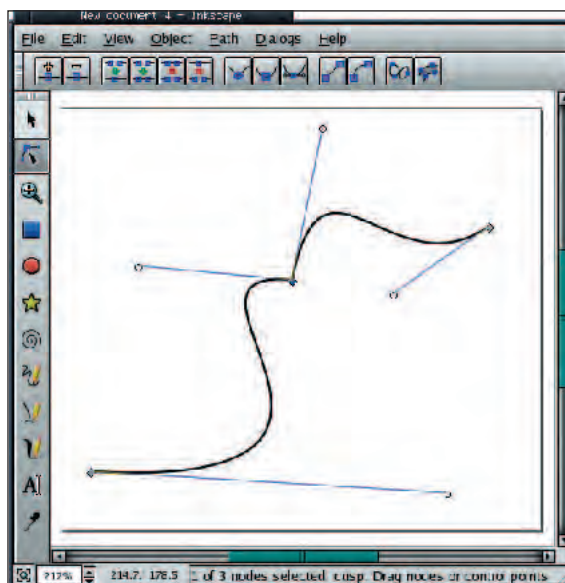
PART I - GET FAMILIAR WITH NODES AND HANDLES

What is a Bezier Curve?

My wife recently started work with *Adobe Illustrator* and was having problems learning how to manipulate paths. Being self-taught, she sometimes comes to me with questions about the tools or terminology. One day, she came in frustrated and said she would never be able to figure out those stupid Beezers. "Beezers?" We still call them that for a laugh.

They're actually pronounced *beh-zee-AY*, after Pierre Etienne Bézier, the French mathematician and engineer who developed this method of computer drawing in the late 1960s while working for Renault, the car manufacturer. Logo design generally starts with simple text or a standard shape such as a box or circle. These objects are then converted to a path, which gives them nodes that can be moved and manipulated to change the shape of the path. A path and its nodes are collectively what make up a Bezier curve.

Inkscape shows nodes as square boxes dispersed along a curve. Two nodes make up a line segment and a line segment can be manipulated by adjusting the handles for the two nodes. Handles are shown as small circles attached to the node by a straight line.



Activating and grouping

A Bezier curve with nodes and handles. Clicking on a node makes it the active node. Holding the shift key down and clicking on another temporarily groups the nodes together. With two nodes selected, you can click on the Curve Segment button to add handles between the two nodes. If the nodes are very close together, zoom in a bit to have a better chance of grabbing the handles instead of a node.

WORKING WITH HANDLES

Corner, Smooth and Symmetric

A node can have a line segment or curve extending out from either side. Each side that has a line segment or curve will have a handle. You can drag the handles around to change the shape of a curve. The behaviour of the curve with respect to the handle depends on the mode of the handle. There are three handle mode buttons in the node button bar: Corner, Smooth and Symmetric. Corner mode allows you to move a node's handles independently of each other. Dragging the two handles toward each other will produce a sharp corner in the curve at the node. This means the right-side handle does not affect the left-side curve.

Smooth also allows you to move the handles independently, but the curve at the node remains smooth. This means the right-side handle can affect the left-side curve in order to keep the curve at the node smooth. Symmetric mode forces the two handles of a node to act in unison. If you experiment with them, you will see that they keep the same distance (but in opposite directions) from the node and are always 180 degrees apart.

There are other buttons in the button bar that you will want to get familiar with later. The most important are the Line Segment and Curve Segment mode buttons. The former will *remove* handles from the selected nodes, while the latter will *add* them. When you create an object and convert it to a path, the nodes initially may not have handles. If they don't, select the nodes and click on the Curve Segment mode button to add handles. The Line Segment mode will remove handles between two nodes and cause the segment between those nodes to become straight.



Handle buttons
From left to right: Corner Mode, Smooth Mode and Symmetric Mode.



Segment mode buttons
Line Segment and Curve Segment mode buttons.



Node button bar

These buttons manage Bezier nodes and handles. If a node already has handles, then they will be displayed when you click on the node. Also, the handles for the nodes to either side of the selected node will display their handles as well. The handles allow a straight-line segment to be made into a curved path. To convert the curve back to a straight-line segment, click on the Line Segment button.



WORKING WITH NODES

Keep it Simplified, stupid!

More often than not, you'll find yourself removing nodes from objects you convert to paths. You can do this quickly with the Path>Simplify menu option, but if you have to do it manually, you can use the Delete Node button. The node buttons work with one or two selected nodes. Selecting more than two nodes will cause these buttons to simply ignore your request.

Simplify is a great tool for converting a scanned object (converted to a path with the AutoTrace Tool,

which is a separate but important tool in itself) to a path with only a few nodes. You can do this with text too. Imagine converting a script font to a path – it will have a jumble of nodes for all the curved segments. Use simplify to reduce the node count and create a personalised font of your own. Try typing "Hello World" in any script font, convert the object to a path, click on the node tool and then run Path>Simplify several times to watch how the font changes.



Managing nodes


From left to right: Add, Delete, Merge two nodes, Join two nodes with a line segment, Split the segment between two nodes, Convert a single node to two separate nodes.

TUTORIAL GIMP & Inkscape

PART II - RECREATING THE LOGO


Deform a square

We start with a simple square. Select 'Path>Object to Path' to convert the object to a path. Click on the node tool, and you'll see nodes at each of the four corners of the square. The square is tall and thin. Click on the left-side ruler and drag out a vertical guideline to the middle of the page. Select the two nodes at the top of the square and drag them to the right, effectively shearing the rectangle. Drag until the top-left node snaps itself to the guide. Now, select just the bottom two nodes and shear to the left. Drag until the bottom-right node snaps to the guide. Remove the guide by dragging it back into the left-side ruler.



Sheared box with guide
The box has been sheared when the nodes snapped to the guide.

Select the two left-side nodes. Click on the Curve Segment button. Adjust the handles as shown. This creates the left-side curve. The amount of adjustment is not easy to measure with the current version of *Inkscape*, so you'll just have to eyeball it for now. Repeat this process for the right-side nodes. Finally, select 'Dialogs > Fill and Stroke' and choose **yellow** for the fill colour of the box. Set the stroke width to **1 pixel**.



Sheared box with curve-adjusted handles
The handles have been adjusted to create the curves on either side of the box.

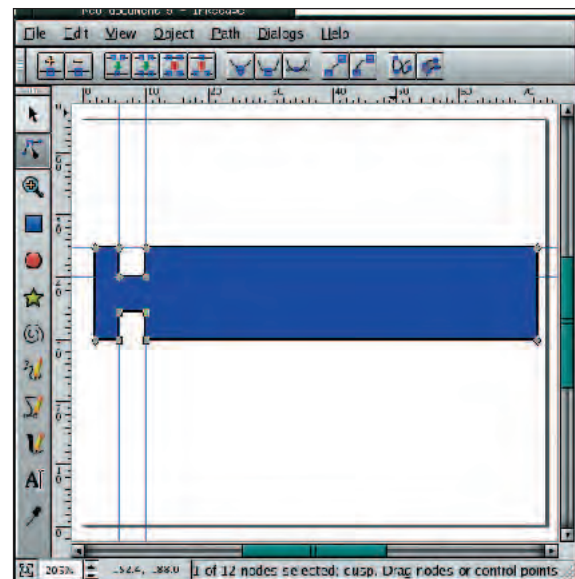
Drawing a pen

Creating the pen requires no handle changes, but does require adding new nodes and repositioning them. Start with a long thin horizontal box – this will be the pen's shaft. Make it moderately large – remember that you can easily resize objects in *Inkscape* without distorting them. Add two vertical guides toward the left side of the box and four horizontal guides evenly spaced with two of the horizontal guides on the outside edges of the box. Convert the box to a path and click on the node tool.

Add two nodes between the top nodes of the box. Drag the two new top nodes to where the vertical guides meet the top horizontal guide. Select the two nodes and add another node between them. Drag this new node down to where the left vertical guide meets the second horizontal guide. Select the new node and its right-side neighbour node. Add a node between these two and drag it to where the right vertical guide meets the second horizontal guide. You now have one side of the stem complete. Repeat this process for the bottom-side of the box to create the other side of the stem.

To the right of the pen shaft, draw another box the same height, but not as long. You can use the horizontal guides for this. Convert the box to a path. Delete the top right node. Move the bottom-right node up so that it is visually centred between the top and bottom horizontal guides. This is the pointed part of the pen.


Using the Bezier tool, draw a straight line from the stem to the right side of the pen shaft. From the 'Dialogs>Fill and Stroke' window, set the width of the stroke to **2 pixels** and the stroke colour to **white**. Draw a similar line from the left side of the pen point to the right tip. These two lines act as reflected light from the edge of the pen.



Pencil with stem nodes and guides

The stem on the pen has four nodes on each side.

Select the Star tool. Click on the pen shaft where the star will go, and drag to create a polygon shape. Hold down the right mouse button on this shape to view an object menu, and choose the Star Properties menu option. Make sure the Flatsides option is set to **False**. Set the number of sides to eight. The object



Reflective star

A reflection is made from a star shape with multiple points and a small internal radius.

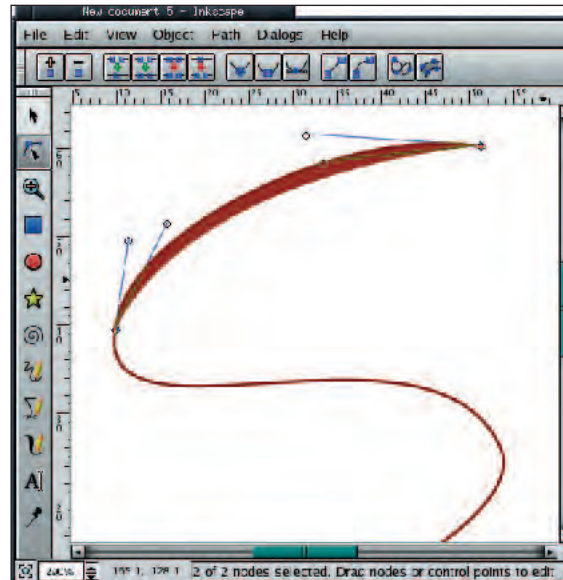
should have two nodes: one for the inner radius, and one for the outer radius (switch to the node tool to see them). Drag the inner radius toward the object centre to create longer, thinner spikes. Change the stroke width to **0** and set the fill colour to **white**.

Finally, select all objects in the pen – including the white reflections and star – and group them as a single object. Resize so the pen is proportionate to the yellow shape. Rotate the object -60 degrees (the amount is shown at the bottom of the window). If you find the size of the pen is not right after the Rotate, undo the rotation and resize again. Note: *always* do the resizing from the horizontal position, not from an angled rotation.

Position the pen group so the pen acts as a 'top-left to bottom-right' stroke of the letter X, with the modified square shape acting as the other stroke.

The swoosh

Draw a reverse-Z shape with the Bezier pen tool. Switch to the node tool. Add handles as needed, and adjust them to make it a smoother S shape. Draw a triangle on the top part of the S, again with the Bezier pen tool. Zoom in to view the nodes' intersection with the S shape. Reposition the nodes to fit on the curve. Adjust the handles for the triangles nodes so that two sides follow the S. Delete the last node of the triangle – we just used it to keep part



The swoosh with triangle nodes

The swoosh has a modified triangle to give it width.

of the path out of the S shape. Adjust the remaining handles so that the segment follows the S closely, but doesn't intersect it. Set the fill for the triangle to match the S shape. Repeat this process for the bottom part of S curve. This creates fat sections on the S curves. Group all these components into a single object.

Text and shadows

The major pieces are complete. All that is left is to add some text and shadows. For the text, set the stroke width to **1** and fill with appropriate colours. Position the text as seen in the original logo.

For the shadows, duplicate the original box we reshaped and filled with yellow. Fill this duplicate with black and offset a small amount to the right and down from the original shape. Then lower the duplicate to the bottom (Object>Lower to Bottom). Repeat this process with the pen shape. After you lower the pen shadow to the bottom, raise it twice with Object>Raise. This puts the pen shadow above the yellow shape, just as in the original logo design.

Your final result may not look exactly like the original logo, but fear not: this is a practice session. The work of learning a new tool predominantly focuses on imitation of existing work. The more you practise to imitate, the better you learn the tool. And that will eventually lead to logo designs of your own... **LXF**

INKSCAPE

www.inkscape.org/

Inkscape's main motivation is to provide the Open Source community with a fully XML, SVG, and CSS2 compliant SVG drawing tool. Additional planned work includes conversion of the codebase from C/GTK to C++/gkmm; emphasising the use of a lightweight core, with powerful features added through an extension mechanism, and establishment of a friendly, open, community-oriented development process.

The current release of *Inkscape*, version 0.38, does not support layers. However, layer features are being added quickly and extensive layer support is expected in the 0.39 release.

Inkscape is available for both Windows and Linux users: download your copy from www.inkscape.org/download.php



NEXT MONTH

If there's any aspect of artwork production under Linux – whether it's using *The GIMP* or other art apps – that you would like to see featured as a tutorial in future issues, or if you have any ideas about particular tasks you would like to see as a step-through, please drop us a line at linuxformat@futurenet.co.uk with 'Art tutorial suggestions' as the subject line, and we'll pick the best for publication!

DESKTOP PUBLISHING

Get creative with Scribus

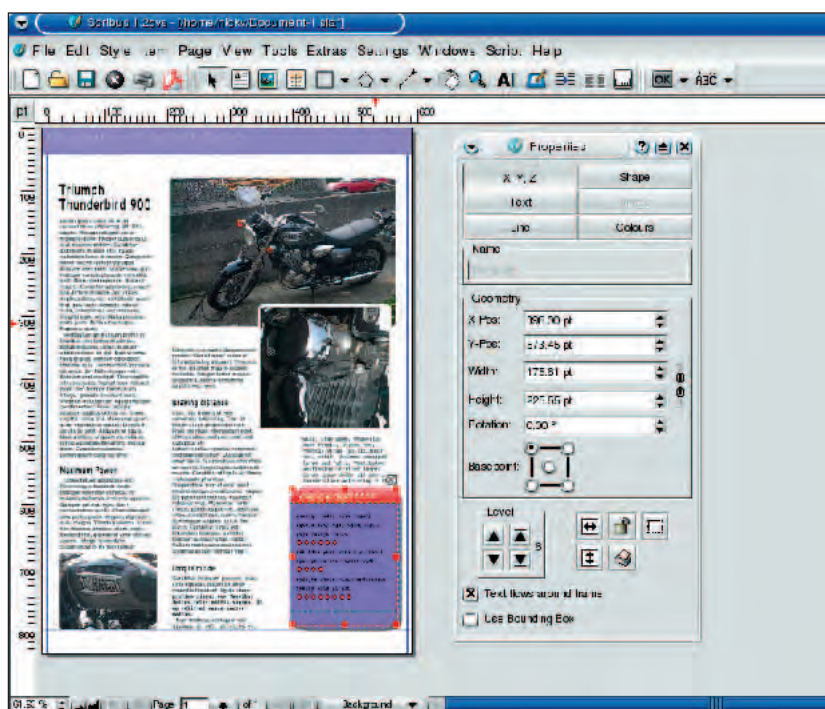
Nick Veitch advocates using Linux's DTP tools more creatively - you probably never realised quite how much you can do with boxes and text...




Since *Scribus* has evolved and become rather good, there are quite a few tutorials available on how to do the basics: set up pages, import graphics and so on. However, a lot of this is pretty obvious if you have already used a DTP package before – *Scribus* is pretty easy-to-use, after all.

So, instead of just showing you how to create a text box and import some text, we're going to be just that little bit more creative by addressing a few less-common tasks, and find out more about the software along the way.

This time, we're going to look at combining objects to make a more complicated graphic device – a colourful box with rounded corners and a shadow that will brighten up your pages. Obviously once you can do this, you can alter the design to best suit whatever publication or DTP project that you are working on.




Scribus is now about as capable as *QuarkXPress v3* – an expensive proprietary DTP app that has been a standard worldwide in magazine and news publishing for a decade.



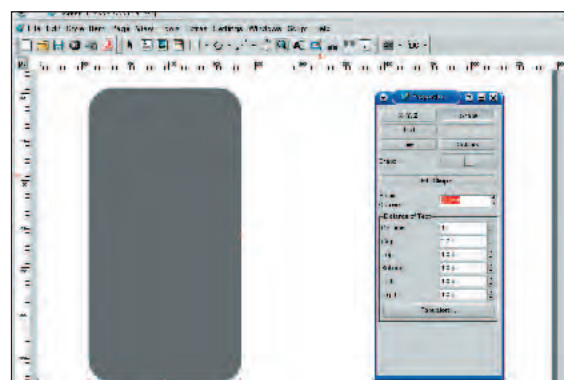
1. Getting started

Create a box. The contents don't matter, but we'll choose a text frame for convenience. Select the Text Frame tool from the menu and draw out the rough size of the box you want. You can tweak it later to fit into any grid designs you might want to fit it into. Now, right-click on the frame object and bring up the properties (or just select Properties from the Tools menu). You can name the object – handy if you think you'll be reusing it at a later date.



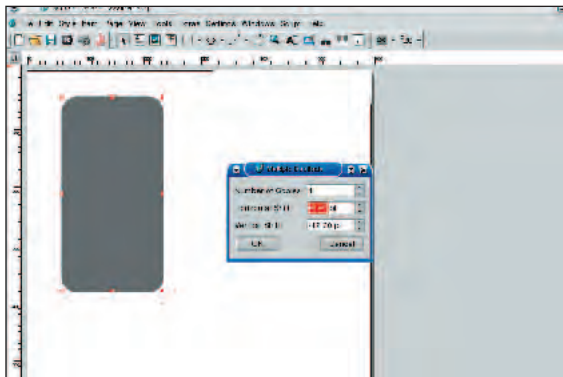
2. Colour, light and shade

The first thing we should do is adjust the colour. This is because it's rather difficult to see the other changes we are going to make if the box doesn't show up against the background. For reasons that will hopefully become clearer later on, we are going to create the shadow first. So, click on the Colours button and select Black, with a Shade of about 50 or 60 per cent for a nice understated shadow colour. Click on Line and make sure the line weight is set to 0.




3. Get in Shape

Click on the Shape button to get to the Shape controls. Leave the box as a square, but in the box marked 'Round corners,' enter a measurement of about 16pt or so. The exact measurement will depend on what you think looks good, the size of the box you drew and so on. Adjust it up and down until it looks right (and if you have a wheel-mouse, you can do this by sliding the wheel up and down while the pointer is over the 'Round Corner' widget).



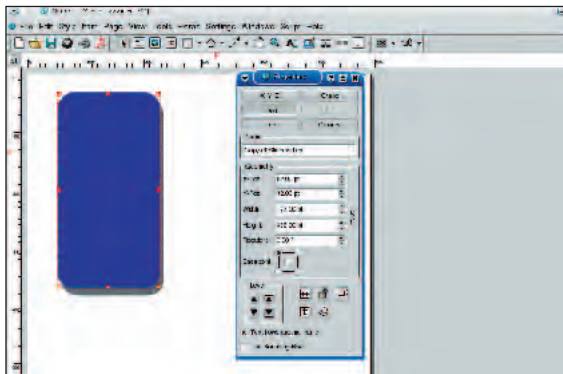
4. Make sure the offset is negative

With the shadow in place, create the box it is a shadow of. Close the properties window, check the text frame is still selected, and choose Item>Multiple duplicate: we only want one duplicate, but this dialog lets you specify an offset; you can have different horizontal and vertical offsets if you like. About half the value of your corner size looks good; our corner here is exaggerated for tutorial purposes.



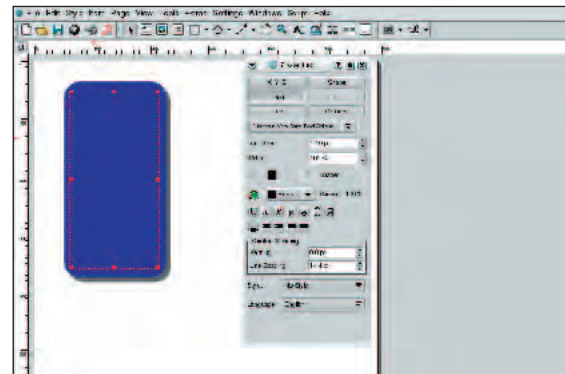
5. The duplicate frame is created on top of our shadow

Now we want to edit the duplicate's properties. Using the pointer tool, click on the top frame and choose the Colours Button again. It's up to you what shade you want to use, but remember that you'll want to run some text over the top of it. We've chosen a nice deep violet colour we created for the purpose (if you haven't got a colour, you can make one now - just choose Edit>Colours from the menu.)



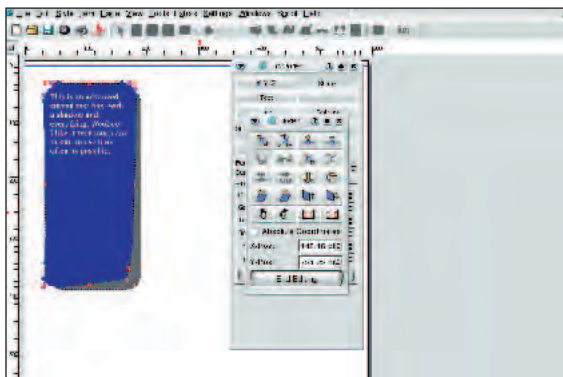
6. Time to make final adjustments

With the colour chosen, you get a good idea of what the finished box looks like. With the top frame still selected, adjust the offset between it and the shadow in either the Properties box, or using the cursor keys. Each keypress makes a positional adjustment. When you are happy, hold down Shift and click in an area of shadow to select it too. Now press Control-G to group the items together.



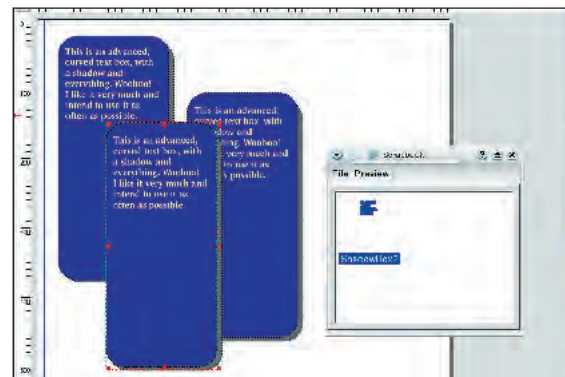
7. It's a frame-up

Although the examples we are using here are text frames, and we do want to put text in them, you may find that the text doesn't start where you would like - most likely, it will start indented in one of the corners. To reposition text, the best work-around is to create another blank text frame and put it on top. Once you have created this frame, you can group it with the others.



8. Node editing

For extra weirdness, you can alter the geometry of the box to make it appear that the upper box is at an angle to the paper. Ungroup the boxes, select the top box and open up the properties. Click on Shape, and then the Edit Shape button to bring up the node editing window. It's fairly simple to add a few nodes and edit them to make it look like the top 'page' is curling outwards. Remember to group them together when you have finished.



9. Finishing touches

Open up the Scrapbook from the Tool menu. With the grouped object selected, just drag-and-drop it onto the Scrapbook page. Now you can re-use this element on any page you design. Remember to save the scrapbook before you quit! There are any number of things you can add to the box (known in publishing as a 'boxout') to make it fit in with the style of what you are trying to do: add a few graphic elements with picture boxes for example.

GETTING SCRIBUS

The Scribus software is available on the LXF coverdiscs this month, but if you want to get hold of the very latest test-build version, and check out some good online tutorials, head over to the newly opened www.scribus.net site.

Game programming

At last, our game becomes playable, as enemies fling themselves to a fiery death at the hands of **Paul Hudson...**



Recently, Nick lent me *Midway Arcade Treasures*, which is a compilation pack of games that were originally released for the arcade but have been ported to the Playstation 2. Some of the games really are classics, such as *Gauntlet*, *Paperboy*, and *Rampage*; others bring back fond memories, such as *Smash TV*, *Toobin'*, and *Road Blasters*; whereas others are pretty obscure, such as *Vindicators* and *Splat*. Still, it was marvellous fun to be playing *Rampage* again after all these years, and it really makes you think how important it is to get the gameplay right before you start focusing on the graphics.

I was chatting with Richard Cobbett just yesterday (PC Plus Features Editor, and large-scale freelance reviewer for all of Future Publishing's gaming titles, both PC and console) and he was bemoaning the lack of good games for Linux. The main thrust of his argument – perhaps quite rightly – is that Linux game programmers spend too much time cloning poor games (see, for example, our review of *Northland* on page 22) and not enough time showing originality. You see, there are a large number of classic games out there that actually deserve being cloned – such as *Syndicate*, for one – and yet all we seem to get are new implementations of a graphical system for *NetHack*. That's not to say these aren't admirable projects, but it's a shame that the only really good game that comes on Linux live CDs is *Tux Racer*, and even that fails to run on most computers because no 3D card driver is bundled!

Though *Trout Wars* is really just a reference implementation of all the ideas I'm trying to teach, I really hope you already have your own game ideas brewing, and perhaps have even made it to the storyboard with some working ideas. The beauty of Open Source is that anyone – even someone who learnt SDL programming just a month ago – can make a game that gives thousands of people happiness and get widespread distribution. For example, the game that the LXF team plays most in the office is *Crack Attack* (<http://aluminumangel.org/attack/index.html>). We spend hours on it, and really it's not complicated, so we hope you're learning from others and remember that it's important to get your first few games out to the public – even if they're less than perfect – as Linux gaming is all about getting help and input from others!

Maybe at the conclusion of the *Trout Wars* series, LXF should run a competition on which reader (or reader team) could make the best game! This issue, we're going to be making enemies appear randomly in the game and also allow the player to shoot at them, which means we've almost got a complete game!

It came from the fish tank

Our first task is to get the enemies spawning randomly on the screen and moving towards the player. This requires exploring a new area of SDL: timers. Go back and look at our call to **SDL_Init()**: **SDL_Init(SDL_INIT_AUDIO|SDL_INIT_VIDEO|SDL_INIT_TIMER)**. We've already looked at **SDL_INIT_AUDIO** (our music) and **SDL_INIT_VIDEO** (our graphics), but now we need to look at what **SDL_INIT_TIMER** gives us. Games need to be able to track time at a very fine resolution, which is why SDL has its own dedicated routines for tracking time. These timers work at the

millisecond resolution and allow you to specify a callback function that will execute after the specified number of milliseconds have passed. Of course, it's important to remember that SDL is platform-independent, and that every platform has different timer granularity. This is further muddled by the fact that loaded systems – such as a Linux box doing various other things at the same time as running *Trout Wars* – will be less likely to hit their expected granularity. The recommended minimum resolution is about 10ms; which means if you ask for 36ms, you should expect about 40ms. This means you can get about 100 callbacks a second, which should be more than enough for every need!

We want to create enemies regularly, so these callbacks are perfect. The first step is to add this line of code to *TroutWars.h*, just below the **std::vector Enemies**:

```
int EnemyDelay;
```

We'll be using this variable to track the time to be used for the callback function. We also need somewhere to store a handle to the callback timer itself – this is important because we may, at some point, want to stop the callback timer from executing, and this requires a handle to the timer. So, add this line beneath **EnemyDelay** in *TroutWars.h*:

```
SDL_TimerID LevelEventTicker;
```

Now, onto the actual task of handling a callback function. If you've ever used them before, you'll know that the defining attribute of callback functions is that they must fit a very precise function prototype. That is, they must accept a set number of parameters of set types and there's nothing you can do to work around this. The function we're going to be using for callbacks is **LevelEventTick()**, so declare it in *TroutWars.h* like this:

```
void LevelEventTick();
```

If you recall from last issue we added the following line to the **CTWGame()** constructor:

```
Enemies.push_back(new CEnemy(this, EnemyTypes[0],
SCREEN_WIDTH - EnemyTypes[0]->sfcPicture->w,
SCREEN_HEIGHT / 2));
```

What that does is add a single enemy to the game, positioned at the right-hand side of the screen centred vertically. We're going to migrate that into our callback function with the difference that the vertical position needs to be randomised. Back in LXF55 we added the inline function **randrange()** to get a quick and easy number randomisation, which is another piece of the puzzle done already. This gives the following initial implementation of **LevelEventTick()**:

```
void CTWGame::LevelEventTick() {
    Enemies.push_back(new CEnemy(this, EnemyTypes[0],
SCREEN_WIDTH - EnemyTypes[0]->sfcPicture->w,
randrange(0, SCREEN_HEIGHT_TRIMMED)));
}
```

Now, here's the prototype for the SDL callback function:

```
static Uint32 StaticLevelEventTick(Uint32 interval, void *param);
```

Clearly, our prototype is different, but what's stopping us from changing that? Well, the sole line of code in the function calls **Enemies.push_back()** and passes in **this** for the current **CTWGame**. The prototype we need to stick to would make the

SUPPORT SDL

SDL is wholly supported by funding from the community, so why donate just a little money to help pay for development costs? The SDL homepage is at www.libsdl.org and there are instructions online for how to donate.

function static – that is, it doesn't have a local reference to itself because it may not have an instantiated object. So, our code cannot be used inside a static function – natch. However, we can cheat to get around this by having a static function that fits the static prototype but goes ahead and calls our non-static function. This can be achieved through the second parameter to **StaticLevelEventTick()** – **void *param**, which is there to allow us to pass in whatever we want. In order to be able to call our running **CTWGame** object we need to pass it in. Here's how it works:

```
uint32 CTWGame::StaticLevelEventTick(uint32 interval, void
*param) {
    CTWGame* thegame = (CTWGame*)param;
    thegame->LevelEventTick();
    return thegame->EnemyDelay;
}
```

You'll need to put the prototype for that into *TroutWars.h*. Note that the **param** parameter is typecasted to **CTWGame**, then we use that to call the real **LevelEventTick()** function from there. The value being returned is key, as if you return the same value as the callback was set with it just resets it so it gets called again. If you return a different value, the timer gets changed to the new value.

To get the timer started, add this code to the end of the **CTWGame()** constructor:

```
LevelEventTicker =
SDL_AddTimer(EnemyDelay,StaticLevelEventTick, this);
```

Try running the game – you should see fish being created along the right-hand edge of the screen as predicted! In the code supplied I've added a call to remove the timer once the game is finished.

Action stations!

Our next goal is to have the fish move towards the player as they are created, and – ahem – it's easy. We already have the **CEnemy::Update()** function sitting there and being called for every enemy as part of our **DrawScene()**, which is the perfect place for our movement code. We already have code in our **LoadEnemies()** function that grabs the speed our enemies should be moving from the *EnemyTypes.txt* file, so we just need to merge the two to make our enemies move.

So, edit your **Update()** function to this:

```
bool CEnemy::Update() {
    this->xpos -= this->EnemyType->speed;
    return true;
}
```

This just shifts each enemy to the left by the number of pixels specified in the *EnemyTypes.txt* file. Is it really that simple? Yup! Go ahead and try it out and watch the fish move neatly across your screen. But wait – there's more! Notice how the fish still start from the right-hand edge of the screen as opposed to *off* the screen, which makes their creation quite obvious. Also notice how the Trout Wars binary slowly chews up resources as it creates more and more **CEnemy** objects without destroying them.

So, to complete the movement of our fish, we need to rewrite it so they start off screen, and destroy themselves once they go off the left-hand side of the screen. If you were wondering why **CEnemy::Update()** returned a Boolean and **CEnemy::Draw()** returned void, you're about to find out!

Change your **Update()** function to this:

```
bool CEnemy::Update() {
    this->xpos -= this->EnemyType->speed;
    if (xpos < -this->EnemyType->sfcPicture->w) {
        return false;
    }
```



Many fish, moving, and being shot at – it's the first time in three issue-worth of tutorials that we've got a vaguely interesting screenshot!

```
}
return true;
}
```

That little snippet will make the function return **false** when the enemy is totally off-screen (including the right-hand edge of its picture), which solves part of the deletion problem. Note the **Update()** returns **false** when the object should be deleted: we just need to implement the actual deletion. As our **Enemies** storage area is a **std::vector**, there are special rules for deletion – if you delete from it while part way through looping through it, you may end up accessing bad memory. Instead, we need to use a special STL algorithm to do the job for us. First up, the easy bit – drawing enemies in your **DrawScene()** function should be done like this:

```
for (unsigned int i = 0; i < Enemies.size(); ++i) {
    if (Enemies[i]->Update()) {
        Enemies[i]->Draw();
    } else {
        delete Enemies[i];
        Enemies[i] = NULL;
    }
}
```

If **Update()** returns false because the enemy is off the screen, we delete the enemy and set its memory to null. This is key, as we'll see in a moment. However, if you try to run the program now it should crash because although the memory is nulled, it is still in the **Enemies** vector and thus will be treated as if it were an active enemy. This is where the algorithm comes in: once all the enemies that should no longer be shown have been set to **null**, we need to go through the vector and remove any items that are null. STL handles that with an algorithm called **remove_if()** that allows you to specify a test to be performed on your vector and have it be sorted so that all those that match the condition are at the end of the vector. The first step is to include the STL algorithms in *TroutWars.h*, like this:

```
#include <algorithm>
```

We also need to define the test we want to run on the vector, which, as we've already set each dead enemy to null, should simply return true if the checked item is null, like this:

```
bool IsNull(void* somepointer) { return (somepointer == NULL); }
```

With that done, we can do the actual deletion of the null elements, with this code:

```
Enemies.erase(remove_if(Enemies.begin(), Enemies.end(),
IsNull), Enemies.end());
```


That's really two lines of code put into one. The first part is this:

```
remove_if(Enemies.begin(), Enemies.end(), IsNull)
```

That goes from **Enemies.begin()** (the start of **Enemies**) to **Enemies.end()** (the end of **Enemies**) taking each **CEnemy** and passing it into our **IsNull** function. If **IsNull** returns **true**, that



TUTORIAL Game Programming



Our hero's spaceship – and your passport to fishy destruction. Ignore the magenta, as that's just being used at this stage to define transparent areas.

particular **CEnemy** gets put to the end of the **Enemies** vector. The remainder get shuffled forwards to fill the gaps, but, crucially, they remain in the same relative order. After running our **Enemies** vector through **remove_if()** and **IsNull()**, the end result is that all the null **CEnemy** objects are at the end of **Enemies**.

The return value of **return_if()** is an STL iterator that points to the first element in **Enemies** that is null. So, everything from there to the end is definitely null, so we pass that directly in to the **erase()** function as the start point to erase, **Enemies.end()** as the end point, which means “erase everything from the first null element to the end”. Perfect.

With the fish destruction complete, the last change to make is to have the fish get created off screen. To make this change, we need to edit the **LevelEventTick()** function – currently it sets the initial X position of our fish as **SCREEN_WIDTH** minus the width of the fish texture, which places it just on the screen. Change this to remove the width of the texture, and it will appear just off the screen – here's the code:

```
void CTWGame::LevelEventTick() {
    Enemies.push_back(new CEnemy(this, EnemyTypes[0],
    SCREEN_WIDTH, randrange(0, SCREEN_HEIGHT_TRIMMED)));
}
```

Compile the game and run it – our two problems should be fixed.

Shooting fish

Now that fish fly randomly towards our player, we almost have a game. Things can be made much more interesting if we allow our player to shoot back at the fish as they come, so this is what we'll be adding next. As with enemies, lasers will be split into types of laser (in the class **CLaserTypes**) and actual in-game lasers (in the class **CLaser**). However, unlike enemies, there will be a fixed number of lasers in the game at any one time, so rather than using a `std::vector` we'll be using a static array.

First things first, let's define the two classes: what should a laser type and a laser know? Here's the code I came up with:

```
class CLaserType {
public:
    int damage;
    int speed;
    CLaserOwnerType type;
    CLaserMovement movement;
    int width;
    int height;
    SDL_Surface* sfcPicture;
    Mix_Chunk* mixSound;

    ~CLaserType();
};

class CLaser {
public:
    int xpos;
    int ypos;
    CLaserType* LaserType;
    bool active;

    CLaser();
};
```

So, a laser type has a damage amount, a speed, a width, a height, a picture, and a sound when fired. A laser has an X pos, a Y pos, a laser type, and knows whether or not it's active. That last

Boolean is quite important because we have a static array of lasers – we'll be creating them all when the game starts up and just setting them to inactive or active depending on what's going on. This is so that we don't badly fragment memory when firing lasers at high speed.

You may have noticed that I skipped over the **CLaserOwnerType** and **CLaserMovement** variables. These are special values that define who owns the laser (is it the player or an enemy?) and how it moves (is it left to right, or right to left?) They may also have further values in the future, so I've defined them as an enumeration (commonly called **enums**) – a set of values that can be typed in like text but are treated like integers for maximum speed. Here are the two **enums**:

```
enum CLaserMovement { ImLTR, ImRTL };
enum CLaserOwnerType { ItPlayer, ItEnemy };
```

You'll need to put the **enums** at the top before all the classes, and put our two new classes at the front of the other classes.

To create our two static arrays for lasers and laser types, add these three lines to your definition of **CTWGame** in **TroutWars.h**:

```
CLaserType* LaserTypes;
CLaser* LaserShots;
int LaserCount;
```

The third line there will store the number of the last-used laser. Without this, each time a laser is fired, we'd need to skim through all the lasers looking for the first inactive one – quite a time-wasting process.

Because we have a fixed number of lasers, we're going to define that as a macro for ease of use, so add this text to **TroutWars.h** beneath the definition for **MUSIC_INGAME**:

```
#define MAX_LASER_SHOTS 127
```

Setting up the lasers needs to be done in two parts. First, the memory needs to be allocated for all our 127 lasers. Second, the surfaces and sounds need to be loaded. To set up the memory for the lasers, add this code in the **CTWGame()** constructor in **TroutWars.cpp**, above the line **Music = new Mix_Music[MUSIC_NUMSONGS];**

```
// pre-create laser shots
LaserShots = new CLaser[MAX_LASER_SHOTS + 1];
LaserCount = 0;

for (int i = 0; i <= MAX_LASER_SHOTS; ++i) {
    LaserShots[i] = new CLaser;
}
```

Loading the surfaces and sounds in is a different function, in the same way we have a **LoadEnemies()** function. Back in **TroutWars.h**, add this line before the definition for **LoadEnemies()**:

```
void LoadLasers();
```

Put the implementation for the function anywhere you like in **TroutWars.cpp**:

```
void CTWGame::LoadLasers() {
    LaserTypes = new CLaserType[1];

    CLaserType* NewLaser = new CLaserType;
    NewLaser->mixSound = Mix_LoadWAV("laser1.wav");
    NewLaser->sfcPicture = LoadImage("laser1.bmp");
    NewLaser->movement = ImLTR;
    NewLaser->type = ItPlayer;
    NewLaser->speed = 20;
    NewLaser->damage = 1;
    NewLaser->height = 7;
    NewLaser->width = 10;
```

OUR IMAGE SOURCE...

Although the source code is GPLed and therefore free for you to use in any shape or form you like, some of images and sounds used in our game were borrowed from *The Games Factory*, a Windows-based game development tool from Clickteam (www.clickteam.com). Clickteam has also created *Multimedia Fusion* and *Jamagic*, which are much more advanced game development programs, both of which are currently being ported to Linux as you read this. Clickteam very generously agreed to place the resources used in this game under the GPL as well, which means you are free to use them as you please, so long as you mention that they are from Clickteam.com.

```
LaserTypes[0] = NewLaser;
}
```

That function essentially hardcodes one type of laser into the game. You're welcome to amend this so that it loads from a file in the same that the enemies loading works, but it's just repetition to print it here. The bitmap and wave sound file are both provided by Clickteam under the GPL, and you'll find them with the source code. The movement has been set to **LTR** ("left-to-right") and the owner has been set to **ItPlayer** so that it's owned by the player. Speed and damage is what you'd imagine them to be, but note that there's a hard-coded width and height as opposed to using the width of the texture. This isn't an accident: by setting these parameters slightly smaller than the actual width and height of the texture, our collision detection code (coming soon!) won't react until the laser has overlapped with the enemies by a small amount, which looks much better on-screen.

There are four more tweaks that we need to make before our laser loading code is complete. First, as we're relying on the **active** variable to tell us what lasers aren't being used, we need to pre-set all the lasers to have **active** as **false**. The easiest way to do this is to have a **CLaser()** constructor, so we will need to add this code to *TroutWars.cpp*:

```
CLaser::CLaser() {
    xpos = 0;
    ypos = 0;
    active = false;
}
```

The second tweak is to have our **CLaserType** objects free their graphics and sounds when they are destroyed, so in order to do this, we need to define a destructor for **CLaserType** in *TroutWars.cpp*, like so:

```
CLaserType::~CLaserType() {
    SDL_FreeSurface(sfcPicture);
    Mix_FreeChunk(mixSound);
}
```

The third tweak is to actually call **LoadLasers()** inside **CTWGame()**, so add this rather obvious line before the call to **LoadEnemies()** in **CTWGame()**:

```
LoadLasers();
```

Finally, we need to destroy the **LaserShots** and **LaserTypes** arrays when the **CTWGame** is destroyed, so to complete the fourth tweak, add these two lines to the **~CTWGame()** destructor:

```
delete[] LaserTypes;
delete[] LaserShots;
```

Ready, aim...

Our lasers now get loaded and destroyed, which means all that remains is to actually fire them. The actual process to do this is quite simple:

- i Take the next available laser, and set it to active
- ii Move the active laser to the current player's position (adjusted for where we want the laser shot to appear)
- iii Set the laser type to what we want (we only have one right now, which makes this easy!)
- iv Move lasercount up one
- v Play a sound

Before we look at the code, I want you to make two changes to your code: first, comment out the text from the background, and second, comment out the random sound playing. These are no longer needed!

To fire a laser, we want a **FireLaser()** function for our **CPlayer** class. It should look like this:

```
void CPlayer::FireLaser() {
    game->LaserShots[game->LaserCount]->active = true;
    game->LaserShots[game->LaserCount]->xpos = xpos + 63;
    game->LaserShots[game->LaserCount]->ypos = ypos + 29;
    game->LaserShots[game->LaserCount]->LaserType = game->LaserTypes[0];
    ++game->LaserCount;
    if (game->LaserCount > MAX_LASER_SHOTS) game->LaserCount = 0;
    Mix_PlayChannel(-1,game->LaserTypes[0]->mixSound,0);
}
```

The **63** and **29** are magic numbers to position the laser at the front of the player's spaceship. The if statement in there is so that our laser usage loops back to **0** if we try to use more than the available number. Other than that, it's self-explanatory – not bad!

Drawing the lasers is also simple – you just need to loop through all the lasers and, if one is active, move it along a little and draw it at its new position. This makes our **CPlayer::Draw()** function look like this:

```
void CPlayer::Draw() {
    game->DrawImage(sfcSpaceship, xpos, ypos);
    for (int i = 0; i < MAX_LASER_SHOTS; ++i) {
        if (game->LaserShots[i]->active == true) {
            game->LaserShots[i]->xpos += game->LaserShots[i]->LaserType->speed;
            if (game->LaserShots[i]->xpos > SCREEN_WIDTH + 100) {
                game->LaserShots[i]->active = false;
            } else {
                game->DrawImage(game->LaserShots[i]->LaserType->sfcPicture, game->LaserShots[i]->xpos, game->LaserShots[i]->ypos);
            }
        }
    }
}
```

As you can see, that moves each laser to the right by the same number of pixels as its speed, checks whether it's off-screen (**100** is arbitrary), and disables it if it is. If the laser remains active,

The final piece in the puzzle is allowing our player to fire a laser using the spacebar.

```
if ( event.type == SDL_KEYDOWN ) {
    if ( event.key.keysym.sym == SDLK_ESCAPE ) { done = 1; }
    if ( event.key.keysym.sym == SDLK_SPACE ) { Player->FireLaser(); }
}
```

Note that the code now first checks that the event is a key down event (as opposed to a key up event or any other type) and then checks whether it was space that was pressed. It's important to do it this way because if you just check whether space is held down it will execute the **FireLaser()** function many times thanks to the speed of the game. I've also added a line in there to handle quitting the game when **Escape** is pressed, which is quite common in SDL games.

As per usual, I suggest you check your code against what I've provided on your CD/DVD, as sometimes I slip in minor changes to help keep the code tidy. *Trout Wars* seems to have become the most popular topic on the *Linux Format* forums recently, so if you have any problems getting your game to work or have any suggestions for future instalments, let me know there! [LXF](#)

COMPILING YOUR GAME

Once you have SDL installed correctly, compiling is quite easy. First, make sure you have both *g++* and *sdl-config* installed on your system – try running them with **--version** and make sure you don't get any errors.

If you have both of these two, compiling is a case of running this:

```
g++ TroutWars.cpp -o TroutWars `sdl-config --libs`
```

That should create an executable called *TroutWars*, which you can run using *./TroutWars*. Note that the quotes around the *sdl-config* part are actually backticks, which means that *bash* will execute *sdl-config --libs* and put its output into *g++*. If you don't have *sdl-config*, you should be able to use

```
-ISDL -ISDL_mixer -ISDL_ttf
```

to get the same results.

NEXT MONTH

We'll be implementing collision detection between the lasers and the fish, and also looking at animating an explosion when fish are destroyed.

INTERACTIVE FICTION PROJECT

Practical PHP programming

This series has been running two years, but only now does **Paul Hudson** *really* start getting practical...

Congratulations! If you've read all the PHP tutorials to date, you've learnt enough about PHP to be able to get out there and feel confident programming pretty much anything. It's now time for us to create a full PHP project in *LXF* – something you can follow through from start to finish, that pushes your skills right to the edge and hopefully beyond, and also gives us something that works really well that we can all be proud of.

Following on from the success of our Trout Wars tutorial, the project we've chosen to produce is an interactive fiction site. If you ever read the old *Choose your own Adventure* (CYOA) books when you were younger, you'll remember you would read a page in the book and at the bottom it would say something like, "if you want to turn left, turn to page 50," and "if you want to fight the ogre, turn to page 59." On the other hand, there are fully fledged role-playing games (RPGs) such as *Dungeons & Dragons*.

Interactive fiction falls into a halfway house between the two; that is, there is a fixed storyline like the CYOA books, but like full RPGs you have a character that can pick up objects and use them, talk to characters, and more. What we're going to construct is a world-creation toolkit that allows writers to create an adventure for others to play. This is potentially a huge project depending on how much control you want to give the adventure writers, so, be prepared for a lot of work! Fortunately, the end result is an exciting system that people can have a lot of fun with.

Best-laid plans

What we're going to produce is big. If you thought Burger King's Extra Double Jumbo Whopper was big, think again: what we're going to produce is *big*. When producing a game system, giving people the flexibility to make what they want out of it is key: hence the popularity of level or scenario editors. While we could make a system in a day, it would almost certainly not be flexible enough to let adventure writers truly bring their games alive.

In order that you fully understand the design of the system, the entirety of this month's tutorial is going to be devoted to the layout and interaction of the tables. *It is absolutely key that you understand how the tables interact, otherwise you'll run into problems later if you get confused.* The focus is always going to be PHP and SQL, so the design presented here may change if I end up spending inordinate thousands of words trying to explain how something works!

To begin with we're going to have seventeen basic tables. Stop crying and relax: they are all useful, important tables that are quite easily understood. *Linux Format's* Art Editor, Julian, has done

a sterling job of putting our complicated table design into one big explanatory picture that's over the page, but I still need to explain what each table and field does. Note that each table has an ID field for its primary key to uniquely identify it in that table and to act as a foreign key in other tables.

Before we get started, you need to know two things: the people who create the adventures will be referred to as *games masters* (GMs), and each character playing an adventure has a set of skills and attributes that defines how good they are at various things. Each character will have a rating for: **Strength, Intellect, Energy, and Agility.**

users This stores information about users who have created accounts on the server. You might want to add more fields, such as name, age, demographics, etc.

■ **Username:** this stores the username of each user. Sounds obvious, we know, but it's important because although users will log in using their email address and password, their name will be shown on the site using the **Username**.

■ **Password:** a standard password field. This won't be encrypted or hashed, although you're welcome to write this functionality simply by calling **sha1()**.

■ **EmailAddress:** this won't be verified at all, but you're welcome to add this. For example, adding an extra field called **Confirm** that was a VARCHAR(40) would allow you to run **sha1()** on their email address + a call to **time()** and have that as their confirmation number. This then gets emailed to their email address, and the user needs to enter that into the site in order to activate their account. So, this should be fine:


```
$confirm = sha1($email + time());
```

■ **DateJoined:** the date their account was created, and isn't strictly required: you will probably find it useful for tracking how popular the site is, though.

guests This stores information about guests on the site. As each person playing the game needs an account, they must either be logged-in users (with an entry in the **users** table) or guests.

■ **CreationTime:** as guests won't be able to log in to their account to continue playing where they left off, storing the time their guest account was created means that we can easily clean up the system of all guest games that were started more than X days ago. As each game could well be quite a bit of information, it's important to clean up like this regularly. >>

Note that ALL fields called "Game" link to the ID field in the **games** table – we tried putting the lines in for all these links, but it just got too messy!



games These are actual adventures created by users for others to play. People must be logged in to create an adventure.

■ **Name:** the name of the adventure, eg *"Tux Strikes Back"*. This is just to give players a way to refer to the game.

■ **Creator:** ID number of the user who created the game.

■ **Difficulty:** how hard the game is. This won't actually affect how the game plays – it's there for the creator to let potential players know what they are letting themselves in for. **1** is *Easy*, **3** is *Medium*, and **5** is *Hard*. There are gaps in there in case we later wish to add something inbetween at a later stage.

■ **DateCreated:** just for information purposes, really.

■ **Info:** a text description of the adventure to whet players' appetites and encourage them to try it out.

■ **MinCharLevel:** each character has a **Level** that defines how powerful they are. By default, characters start out at level **1**, but by setting a higher number in here, GMs can force new characters to at least be (for example) level **10**, and thus more powerful. This in turn means they can make the adventure harder.

■ **PointsAdjust:** each character has a set number of character points to spend adjusting their basic attributes, and GMs can either increase those points to make characters better or decrease them to make them worse.

■ **RandomEncounters:** percentage chance of a random encounter each time a player changes location. Random encounters are unscripted attacks on players that the engine will automatically create to keep players on their toes. Setting this to **0** disables random encounters.

■ **RandomEncounterStrength:** this defines how strong enemies should be when they are met in a random encounter. **1** is *Easy* (should be very easy), **3** is *Medium* (should hurt the player a little), **5** is *Hard* (player should survive, but only just), **7** is *Hardest* (equal in strength to the player), **9** is *Impossible* (well, more or less – greater in strength than the player).

■ **Status:** the status of the adventure. **1** is *Just created*, **2** is *Playable*, **3** is *Down temporarily*, **4** is *Down long term*. Only status **2** allows others to play the game. **3** and **4** are there so that the GMs can make changes to the game without causing problems for players.

■ **Length:** like **Difficulty**, this is just an indication from the GM to the players. **1** will be shortest, **9** will be longest.

characters This is where characters are stored. Note that characters are separate from users, as each user (each physical member of the site) can have many characters (fictional players in games). Also note that each character can only be in one game at one time.

■ **Owner:** ID number of the user that owns this character.

■ **Name:** name of the character, eg *Bork the Indestructible*.

■ **Age:** age of the character. For information only.

■ **Sex:** sex of the character. For info only. **1** is *Male*, **2** is *Female*.

■ **Race:** race of the character. This can affect the min and max ages of the character and may also affect their attributes.

■ **Class:** character class of this character, eg **Fighter** or **Mage**. This affects their attributes as well as their potential for spell casting and healing.

■ **Level:** experience level of this character. This determines how powerful the character is.

■ **ExperiencePoints:** each time a character kills a creature or performs an action, they might receive some experience points. Experience points determine the level of the player.

■ **Info:** free text for the player to describe their character. For information only – this has no bearing on the game.

■ **Game:** ID of the game this character is currently in.

■ **Location:** ID of the room this character is currently in. More on this soon.

■ **GameState:** ID of the game state this character is currently in. More on this soon, although **-1** means "This player has yet to start playing their game."

■ **Strength, Intellect, Energy, Agility:** the four basic attributes that determine how good a character is at doing various things. Increased strength – more dangerous when fighting; increased intellect – they can learn more; increased energy – they can take more damage, and increased agility – they are more likely to hit people in combat and dodge enemy blows.

■ **Vitality, VitalityMax:** the maximum amount of health this character can have is defined by **VitalityMax**, whereas **Vitality** is the current level of energy this character has.

■ **Mana, ManaMax:** the actual amount of spellcasting power the character currently has is **Mana**; **ManaMax** is the maximum amount of mana (spellcasting power) this character can have; .

■ **Hunger, HungerMax:** the latter is the maximum amount of time this character can survive before starving, whereas the former is how close the player is to starvation currently.

races GMs can define the races they want players to be able to choose for their character in a given adventure. For example, if we created an adventure based around Lord of the Rings we might use the races *Hobbit*, *Elf*, *Dwarf*, *Human*, *Ent*, etc.

■ **Game:** ID of the game these races are for. Note that a **-1** in this field means "this race is available in all games that don't specify their own races" – we'll be providing a stock selection of races that are there by default.

■ **Name:** name of the race

■ **Info:** free-text description of the race to explain some history and back-story to players.

■ **PointsAdjust:** GMs can have this race grant a character point bonus (or penalty) to players.

■ **StrengthAdjust, IntellectAdjust, EnergyAdjust, AgilityAdjust, VitalityMaxAdjust, ManaMaxAdjust, HungerMaxAdjust:** these allow the GM to provide bonuses or penalties to these character attributes for this class.

■ **MinAge, MaxAge:** what age characters of this class can be.

classes These are essentially vocations or "jobs" that define what players do best. They can be as generic (eg *Fighter*) or as specific (eg *Berserker*, *Knight*, *Stormtrooper*) as GMs want.

■ **Game:** ID number of the game this class belongs to. Like races, if this is set to **-1** (which GMs will not be able to do) this class will be available in all games where no special classes have been defined.

■ **Name:** the name of this class.

■ **Info:** text description of the class to make it more alive to players.

■ **StrengthMin, IntellectMin, EnergyMin, AgilityMin:** minimum attributes that characters must meet in order to be a part of this class. Average scores for the attributes are **10**.

■ **SpellCasting:** how strong characters of this class can be at spellcasting. **0** is nothing (they will never have any skill), **1** is *Very weak*, **2** is *Poor*, **3** is *Average*, **4** is *Good*, **5** is *Powerful*, **6** is *Master mage*, **7** is *Legendary magician*.

■ **Healing:** how strong characters of this class can be at healing. **0** is *nothing*, **1** is *Very weak*, **2** is *Poor*, **3** is *Average*, **4** is *Good*, **5** is *Powerful*, **6** is *Master healer* and **7** is *Legendary healer*.

■ **FightingAdjust:** this allows GMs to directly affect how good



TUTORIAL PHP

characters of this class will be at fighting. So, while a *Magician* class might have a **4** in **SpellCasting** it might also have a **-3** in **FightingAdjust**, whereas a *Fighter* class might have a **0** in **SpellCasting** and **Healing** but a **+10** in **FightingAdjust**.

■ **ExperienceLevels:** this defines how easy it should be for people of this class to go up levels. As levels are decided by how many experience points a character has, this figure is applied as a percentage of a basic experience point table. That is, a setting of **0** means a character goes up levels at the standard rate, whereas a setting of **100** means the points required to go up levels are doubled, making progress much slower.

gamestates This defines 'chapters' in an adventure. For example, in the game *Monkey Island* (the original) you had to complete *The Three Trials*: find the idol in the Governor's mansion, find the treasure (although it's just a T-Shirt [apologies if you've never played the game and we've just spoiled the fun!]), and defeat the sword master. As you complete each one, the game changes state – you get a cut-scene where more information is presented about things happening elsewhere, people say different things to you etc. That's what our game states are: when the game state is changed, players are given some text describing what's changed, and the game can react differently.

■ **Game:** the game ID that this game state belongs to.

■ **Name:** the name of the game state (players don't see this).

■ **Info:** free-text description shown to the player when they enter this game state.

Please note: each game needs at least one game state called **START**. This is activated when players first start the adventure.

items This stores all the items that can be in an adventure. It does not define where the items are in a live game currently being played by a character – this is covered later.

■ **Game:** the ID number of the game this item belongs to.

■ **Name:** full name of the item, eg *Sword of Doom*.

■ **ShortDesc:** vague description of this item, eg 'Rusty sword'. This is used when the item is on the ground.

■ **Type:** what this item actually is. **1** is *Weapon*, **2** is *Shield*, **3** is *Helmet*, **4** is *Armour*, **5** is *Ring*, **6** is *Scroll*, **7** is *Potion*, **8** is *Food*, and **50** is *Other*.

■ **DamageMin, DamageMax:** how much damage this item can do when used. If a weapon, this dictates how much it hurts the enemy. If a potion, this dictates how much it heals the wearer. If a *Ring*, this dictates how much it hurts enemies. If *Food*, a negative figure dictates how much hunger it removes.

■ **Info:** free-text description of the object. This is used when the item is examined by the player.

■ **UseWrong:** text to be printed when the object is used in a place where it has no use, eg "You can't use your scroll here."

■ **UseRight:** Text to be printed when the item is used in the place where it is supposed to be used, eg "You throw the rock through the window, leaving a large hole you can climb through."

■ **UseRightTrigger:** what to do when the object is used correctly. Triggers are discussed soon.

■ **ExperienceUseRight:** how many experience points should be awarded when the item is used properly.

■ **DeleteOnUse:** set to **1** if this item be removed from the game when used (eg a potion is drunk and thus cannot be used again)

■ **DropOnUse:** set to **1** if this item should be dropped by the player when used; eg 'using' a rock can mean throwing it, so it gets dropped but not destroyed – the player can fetch it again.

■ **LocationFound:** where the item is found.

■ **LocationUsed:** where the item is used. Set to **-1** if this item can be used anywhere.

■ **SellWorth:** how much the thing is worth when sold. **0** means it cannot be sold.

■ **FightingAdjust:** when the item is fought with, this affects how well the player fights. For example, a magical sword might make its bearer a better fighter.

■ **DefenceAdjust:** if this is greater than **0**, this item adds protection to the character that makes them safer against enemies.

■ **MinLevel:** allows GMs to stop powerful items being used by lower-level characters. They can still hold the item, they just can't use it.

■ **CreateWhenState:** allows GMs to set up items to be created when a certain game state is set. For example, once the character returns home after killing the ghost from the previous example, they might find a note stuck to their front door by a dagger that wasn't there before.

■ **AutoStart:** if set to **1**, this item is automatically given to players who play this adventure as soon as they start. This is helpful to give players basic kit such as a bit of armour, a weapon, or, in more imaginative cases, a medallion left to them by the father they never knew that, if the players can just complete some adventure, will reveal all sorts of information...

itemslive This stores where items in each game actually are. For a given game, there is only one entry in the **items** table for it, but as there can be many players playing the game at the same time we need this table also.

■ **Item:** ID in the **items** table for this item.

■ **LocationCurrent:** where this item is (locations are explained shortly), or **-1** if the character has picked it up.

■ **CharacterOwner:** ID in the character table for who owns this item. Note that 'owns' is not the same as 'picked up'. For example, the point of the adventure might be to collect the Amulet of Yendor, which is located thousands of light years away in a distant solar system and will take years of adventuring to find. When the player starts playing, the amulet is copied from the **items** table into **itemslive**, given their character ID as the **CharacterOwner** so that it cannot be seen or touched by other players, and placed in the distant location. This can then be picked up once the player finally reaches their destination.

rooms A 'room' is the term we'll be using for a location in the adventure. This can be as general as 'The forest' or as specific as 'The dark corner behind the cooker in the kitchen' – it's down to the GMs how they plan out their rooms. Each room can link to other rooms so that players can move to and fro freely, although note that a link is not bidirectional. That is, if room A links to room B, room B does not necessarily link back to room A. This is helpful, because the link might be players going down a trap door that closes, or, in more advanced adventures, GMs might create a clone of room A that's subtly different from the original (we'll call the clone room C) so that players can move from room A to room B, see that they can go back to room A (it's actually room C, but it looks like room A), and thus fall into a trap if they go back. Note that the room that is created first is where players start the adventure.

■ **Game:** ID of the game this room belongs to.

■ **Name:** name of the room. This appears at the top of the page so that players know where they are.

■ **Info:** free-text description of the room that describes the

surroundings. This won't include information about what items or people are in this room, as this is handled automatically. Note that GMs will be able to use a simple programming language in this text to make the room more dynamic. More on that soon.

- **SafeToRest:** set to **1** if characters can sleep here. Sometimes denying people the ability to sleep (if the place is too noisy, busy or too dangerous) can be a helpful plot device, as sleeping regains health.
- **SafeFromEncounters:** if the **RandomEncounters** game per cent is greater than **0**, random encounters can happen. However, if this value is set to **1**, characters moving into this room cannot be hit by a random encounter because the room is considered 'safe'.
- **SafeToAttack:** if this is set to **1**, the player may not attack any of the AI players in this room.
- **CallTrigger:** trigger to call each time this room is entered. Triggers are discussed soon.
- **CallTriggerOnce:** trigger to call the very first time this room is entered and never again.

links This table defines the connections between rooms.

- **FromRoom, ToRoom:** IDs in the **rooms** table that defines the link. For example, **FromRoom** of **3** and **ToRoom** of **10** means that a player can move from room **3** to room **10**. As noted already, this does not mean that a player can move from room **10** to room **3** – a separate row needs to be entered for that.
- **ConditionType, ConditionVar:** this link can be shown only when a certain condition is satisfied (more on this soon).

conditiontypes This table defines the various conditions for use in the **links** table. For example, a condition might be "Is a variable set?". If a GM uses this condition for the **ConditionType** in a link, they can specify the name of the variable to check for the **ConditionVar**. Variables are discussed soon! (yes, this might be confusing by now!) The **ConditionTypes** are pre-defined and cannot be altered by GMs. This condition-based room linking is helpful, because it allows GMs to set up situations where (for example) players can only enter a room if they have already done something else, such as pick up an item.

- **Name:** name of the condition to check. This is never seen by anyone but the GM, and only then really as a mnemonic for the ID number because the ID number is faster for the database to work with.

npcs This stands for 'Non-Player Characters', and this is a pretty distinct subsection of the engine that will require work once we get close enough to it! For now, the table is quite simple.

- **Game:** ID of the game this NPC belongs to.
- **Name:** name of the NPC, *eg Jessica Atreides*.
- **Info:** description of the NPC to be printed out when they are examined by the player.
- **Sex:** **1** is *Male*, **2** is *Female*.
- **Race:** ID in the **races** table for the race of this NPC.
- **Class:** ID in the **classes** table for the class of this NPC.
- **Level:** how powerful this NPC is.
- **LocationStart:** where this NPC is

events At last, this is what really makes the game interesting. An event is something that affects the state of the current game or the character. For example, it might set a variable, take away health from the character, or teleport them to another room.

Events cannot be called directly – they are bundled into groups called triggers, and it is these triggers that are called.

- **Trigger:** ID of the trigger in the **triggers** table that this event belongs to.
- **EventType:** ID in the **eventtypes** table that defines what this event does.
- **Variable:** some events need an extra data field to make their change. This is discussed shortly.

eventtypes This stores the various possible event types in a convenient place. My initial implementation of the code has various values hardcoded, as opposed to in a special table like these event types. This will be something we'll be solving *en route*!

- **Name:** name of the event type. This is just for reference during adventure creation. Event types will include SET VARIABLE to set a variable, UNSET VARIABLE to do the opposite, TOGGLE VARIABLE to set a variable if it is unset and unset it if it is set, ADD HEALTH and SUBTRACT HEALTH to heal or hurt the player, and TELEPORT PLAYER to move the character to a new room in the game. Each of these require a Variable to be entered in the **events** table in order to make sense. For example, SET VARIABLE needs the name of the variable to be set (variables don't have values, they are either set or unset), ADD HEALTH needs to know how much health to add, etc.

triggers This table stores the actual triggers that can be called from the game. These are just collections of events and have no inherent magic.

- **Game:** the ID of the game this trigger belongs to.
- **Name:** name of the trigger for internal use by the GM.

variables This table stores variables that have been set by the game. For example, if a guard outside of a palace has been put to sleep by a sleeping draught, a variable such as PALACE_GUARD_ASLEEP should be put in here so that the game knows what to show. Note that variables starting with **_** are not allowed, as this table is also used to handle the **CallTriggerOnce** field of the **rooms** table – when a room has a trigger that should only be called once, a variable is entered into this table called **__IF_ROOM_ENTERED_<ROOM ID>** so that it won't be called again.

- **CharacterID:** ID in the **characters** table for the character that owns this variable.
- **Variable:** name of variable that is set.

Where's the code?

As was said at the beginning, the whole purpose of this instalment has been to detail the various tables that our game engine will use. Please pay close attention to the diagram of how the tables interact – cut it out and keep it if it will help you – because that explains 90 per cent of how the system works. Some things, particularly NPCs, fighting, magic, etc, are yet to be planned out. Fortunately, these don't really affect how the actual adventure creation works, so we're fine for quite some time yet!

As this is such a big project, there are almost certainly various areas that the system can be improved. You're welcome to write in with suggestions if there's something you'd like to see implemented, or if you'd just like to know how it *could* be implemented in your own code. Also, please keep in mind that this is a very complex system we're going to produce: don't worry if you get confused here and there; that's why we have forums on our website! [LXF](#)

NEXT MONTH

Next issue, we're going to be looking at how the tables work together to produce an adventure, and start putting together the website. If there's space I'll also discuss the adventure we'll be hand-coding in SQL so that we can test out the playing engine before we start the adventure creation system.

USING KDEVELOP

KDE Development

PART 4 Jono Bacon uses this opportunity to give us all the right signals about interactivity – there's no point developing a program if people can't interact with it.



As we blaze a trail through the forests of ignorance and gingerly negotiate the swamps of despair on our quest for KDE enlightenment, we should remember that we have actually achieved a great deal in the first three tutorials of this series. First, we genuflected at the temple of *KDevelop*; second, we generated a window of youthful potential; third, we filled such a window with blessed widgets; and this month we stand before yet another mythical trial... OK, so maybe KDE development isn't actually quite as chivalrous and spiritually enlightening as all that, but it is cool, and this month we are going to take our knowledge one step further – interaction.

Interaction is the most fundamental part of modern software. If you take away all of the widgets, wizards, icons and graphics, you will be left with interaction as software's prime asset. Due to the importance of such interaction, we face a challenge in creating an application that is not only responsive, but also intuitive and predictable to the user.

This issue, we are going to explore the Qt and KDE methods of dealing with interaction. Due to the scope of the subject, we are only going to be exploring the fundamentals of this system in this issue, and in future episodes we will expand and improve on what we have learned here. The reason for this is to ensure that we cover the foundations clearly enough to make our more advanced code easier to digest later in the series.

Slotting into place

Due to the Qt foundation of the KDE desktop, both KDE and Qt applications use the same system for dealing with interaction in applications. This system deals with the concept that widgets

essentially have two functions: they are there to be used for whatever the widget was designed for (eg a text box for entering text); and they are also there to indicate when something has happened to the widget.

As an example, if you click on a button in an application with a Graphical User Interface (GUI), you would normally expect that button to do something. If we consider the situation a little more carefully, that button has a few different methods that the user can interact with it: the user might click it (so the button pushes down and then up), press it in (so the button is held down) or the button may even be toggle-able (so the user presses the button and the button stays sunken). Each of these different types of interaction may need to do something different. An example of this is a typical web browser back button – if you press it, the page goes back, but if you hold the button in, you often get a list of previously visited pages.

KDEVELOP TOP TIPS

Boost your developmental skills

Since this series kicked off a few months ago, we have received a number of emails suggesting areas to cover and topics to explore. A number of these emails were really interested in seeing more content about using *KDevelop* for a wide variety of different things. To bring this kind of content into the series, we will be featuring some regular *TOP TIP* boxes with some useful tips and tricks to make *KDevelop* work that little bit better for you. In addition to providing some of these tricks and hacks, I am also eager to hear your own tips: feel free to send them to jono@jonobacon.org.

TUTORIAL KDE Development

TOP TIP

Do you find that the output in the Messages window is a little too much or too little? The level of output is configurable by right-clicking the Messages pane and selecting 'Very Short', 'Short' or 'Full Compiler Output'. Try playing with the different settings to see which offers the level of information you need.

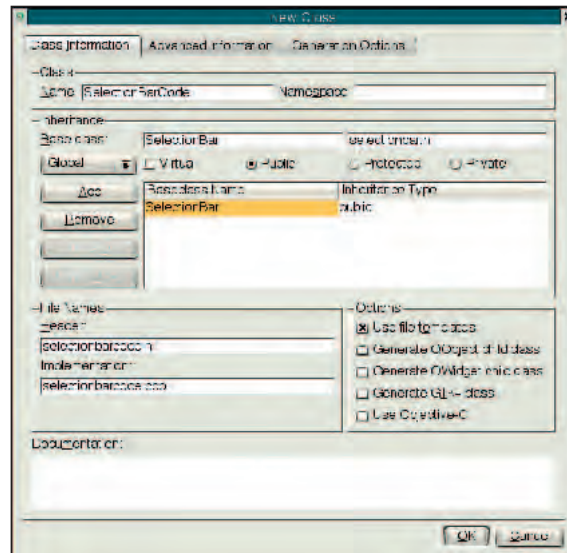


Fig4 Creating the our implementation class.

◀ make a connection with. Select the buttAddAlbum widget. The next column is the Signal selection, and if you click on **<No Signal>**, a drop-down box will allow you to select the relevant signal. We want to allow the user to click the button (a click being the press and release of a mouse button), so we need to select the **clicked()** signal. The next column is the Receiver. This column allows you to select where the slot will be stored. In our case we want the slot to be stored in the SelectionBar class, so select this.

The final column is for the slot name. At the moment we have not actually created any slots, so we need to click on the 'Edit Slots' button. When the dialog pops up you will need to click on the 'Add Function' button to generate a new slot. You should at this point edit the 'Function' text box to change the name of the slot to **addAlbum()**.

All of the default settings in the dialog box will be fine; we are not returning anything from the function, so we can leave the return type as **void** (this means the function does not send anything back to its caller). With the settings complete, click on the 'OK' button and you will return to the 'Connections' dialog box. You can now select your new slot from the Slot column. Click on 'OK' and our signal/slot connection is complete.

All that is left for us to do now is to save the file in Qt Designer and recompile the project. If you keep a beady eye on the Messages pane in KDevelop, you should see that the selectionbar.cpp and selectionbar.h files will be regenerated. When the project is recompiled, you should run the application and click on the new button that we have added. Pay attention to KDevelop while the program is running; when you click on the button you should see this appear in the Application pane:

```
SelectionBar::addAlbum(): Not implemented yet
```

This text will appear whenever you click on the button. This text (**Not implemented yet**) is actually outputted from our generated slot that is within the **SelectionBar** class. This shows that the slot is fully working – whenever we click the button, the default slot outputs that the slot is not implemented yet. We now need to actually implement the slot.

Before we actually implement our slot, we will have a quick look at the code that was generated for us. If we first take at selectionbar.h we can see the following addition:

```
public slots:
    virtual void addAlbum();
```

This code is added to declare the fact that we will be creating an **addAlbum()** function that is a publicly accessible slot. The **slots** keyword is actually a special keyword that is part of the Qt toolkit. When the files are compiled, a special tool called *moc* is run on the code to convert the **slots** keyword into native C++ code that performs some low-level magic to make the signal/slots system work.

The other generated file is selectionbar.cpp, and the following code is added:

```
void SelectionBar::addAlbum()
{
    qWarning( "SelectionBar::addAlbum(): Not implemented yet" );
}
```

Here we can see the actual **slot** function. The code just looks like any other function, and this is because it *is* just a normal function. It is always important to remember that a **slot** is a function that can be used as a slot or as a normal function. Inside our function definition, the **qWarning()** function is used to write to the console that our slot is not implemented; the same text we saw earlier.

Writing our own slots

With our generated code working well, we now need to write some code to create an **addAlbum()** slot that we can edit and not lose when the Qt Designer is changed. To do this we will need to 're-implement' our **slots**.

Re-implementing code is a special feature in C++ that allows us to create a special class that inherits from our generated class so that we can write our slots and they will not be overwritten. To do this we need to create a new class that will inherit from **SelectionBar** and we will then add our constructor, destructor and slot declarations/definitions.

To create the class, select Project>New Class and call the new class **SelectionBarCode**. You should also click on the 'Add' button in the Inheritance section. You need to add **SelectionBar** as a base class. Finally, click on 'OK' and generate the files. When the files have been generated, we will need to add some extra bits of code to add our slot. We will first edit selectionbarcode.h, and your code should resemble this:

```
#include <selectionbar.h>

class SelectionBarCode : public SelectionBar
{
    Q_OBJECT

public:
    SelectionBarCode(QWidget *parent=0, const char
        *name=0);
    ~SelectionBarCode();

public slots:
    virtual void addAlbum();
};
```

The main change to code is in changing the constructor to add the familiar arguments that we added in the previous issues, and also adding the **Q_OBJECT** macro and our public slots block. You may remember that the **Q_OBJECT** macro is added to allow us to make use of the signal/slot system.

Our next file to edit is the selectionbarcode.cpp file. Here is

TOP TIP

If you right click in the main editor on a variable you can view some context options to view different parts of the code where the variable is used. This includes variable declarations and definitions. Some additional options inside this context menu are also well worth exploring.

the relevant code:

```
#include "selectionbarcode.h"

#include <qmessagebox.h>

SelectionBarCode::SelectionBarCode(QWidget *parent, const
char *name) : SelectionBar(parent, name)
{
}

SelectionBarCode::~SelectionBarCode()
{
}

void SelectionBarCode::addAlbum()
{
    QMessageBox::information( this, "LXFGallery",
    "This feature is not implemented yet.\n"
    "We will be developing this feature in a future issue of
Linux Format" );
}
```

Again, we need to add our arguments to the constructor and the base class, and we also add our **addAlbum()** slot here. Inside this slot we use a special class called **QMessageBox** that simply pops up a little dialog box. I have used this class so that we can clearly see when the slot is getting triggered. The **QMessageBox** class has a number of different types of message box, and here we are using the information box. The different arguments (apart from the first which sets the parent) set the the text that is displayed in the box.

At the moment our re-implemented code is complete, but if you compile the project you will still get our generated slot being used. This is because we are still using the **SelectionBar** class to create our interface as opposed to our new **SelectionBarCode** class. To change this we need to edit a few files. First, load up *lxfgallery.h* and change the **#include** that includes *selectionbar.h* to include *selectionbarcode.h*. You will also need to change:

```
SelectionBar * selectionBar;
```

to

```
SelectionBarCode * selectionBar;
```

The final file to edit is *lxfgallery.cpp*. You should change:

```
selectionBar = new SelectionBar(splitter);
```

to

```
selectionBar = new SelectionBarCode(splitter);
```

If you now recompile the project and click on the button, a snazzy new dialog box should pop up. We now know that our slot is working well.

Summing up

In this issue we have explored how to create signals and slots, and we have also explored how we create our signals and slots in *Qt Designer* and then re-implement them later in *KDevelop*. This system may seem a little odd at first, but there is no doubt that the *Qt Designer* concept has been designed to be as flexible as possible. *Qt Designer* would be worthless if you could only design your interface once, and then you could edit it again as you would overwrite the generated code. The current system allows you to edit an interface in *Qt Designer* whenever you like and still be able to retain your slots.

With KDE development, the main concept is to keep the

SOFTWARE REQUIREMENTS

Properly equip your PC for KDE development

These are the specific software requirements for *KDevelop* that you need to follow this series:

- KDE (3.0.2+) www.kde.org
- Qt (3.0.5+) www.trolltech.com/products/qt
- g++ (2.95.3 + or compatible) gcc.gnu.org
- GNU make www.gnu.org/software/make
- Perl (5.004+) www.perl.com
- autoconf (2.52+) www.gnu.org/software/autoconf
- automake (1.6+) www.gnu.org/software/automake
- Flex (2.5.4+) www.gnu.org/software/flex
- Berkley DB (3.0-4.1) www.sleepycat.com

In addition to these essential requirements above, the following tools are optional:

- ht://Dig (3.1.6+) www.htdig.org
A library used for searching for text.
- Valgrind <http://valgrind.kde.org/>
A tool for finding and fixing memory leaks. Memory leaks are important to avoid when create KDE applications in C++.
- GDB (5.0+) www.gnu.org/software/gdb
A full featured debugger for finding and squashing those bugs. *KDevelop* integrates *GDB* within the *KDevelop* environment.

- CVS (1.10.6+) www.cvshome.org

The most popular source control system, *CVS* is commonly used when a number of developers are working on one codebase.

- Perforce (2003.1 +) www.perforce.com/perforce/products.html
Version control and configuration management.
- CTags (5.x +) <http://ctags.sourceforge.net>
Special tags for easing the editing and development of code. This is built right into the *KDevelop* editor.
- Doxygen (1.3.4 +) www.doxygen.org
Automatic generation of programming APIs and documentation.
- dot (1.8.7 +) www.graphviz.org


Generation of graphical class diagrams. This is a highly recommended tool, as you will find that the class viewer is very useful when programming KDE applications.

Most of these tools are available in RPM and Debian package formats, and source code tarballs. If you are using a distribution that allows you to download packages from the Internet and install them automatically, you should have few problems with installing *KDevelop* and its optional extras.

Both CD and DVD readers will find this whole selection on this month's coverdiscs. The specific software from the first part of the list will be included in some later issues, space permitting.

interface design in *Qt Designer* and the code in *KDevelop*. There is no doubt that you could create the same interface in raw code, but this is a time consuming and error-prone process. Although the bulk of our interface design will take place in *Qt Designer*, we will still use code to generate interface elements in parts of our project. As an example, when a user selects a photo album, we will use code to populate our *KListBox* instances with thumbnail images of the pictures in the album. We could not do this in *Qt Designer* as we could not hard code a specific photo album – we need to allow the user to select the album of their choice.

We have covered a lot of ground this month, and we have approached some complex subjects. It is highly recommended that you work through this part of the series a few times and try to understand as much code as you can. We also recommend that you have a read on some aspects of C++ that may seem a little alien to you. This includes virtual functions and polymorphism. There are a number of good books available that teach C++ and a good start is C++ FAQs (there is also a website with the C++ FAQs at www.parashift.com/c++-faq-lite/). If you can familiarise yourself with some of the C++ concepts that seem unusual to you, you will find that the future parts of the series will become much easier to understand.

Some parts of this issue may seem confusing, but don't worry. As the series continues and we explore more areas and write more code, many of these concepts will become clearer. 

TOP TIP

With any kind of technology, separation is often a good thing. In the same way you separate data into a database and web pages into files, it is good to separate out interface creation into *Qt Designer* and functionality into code. What is often the most difficult concept to grasp when learning a new piece of technology such as KDE development is how everything separates out this way.

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! Don't forget there's also advice available from our website – note that registering and posting coherently to the appropriate forum is likely to get you better help, though!

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



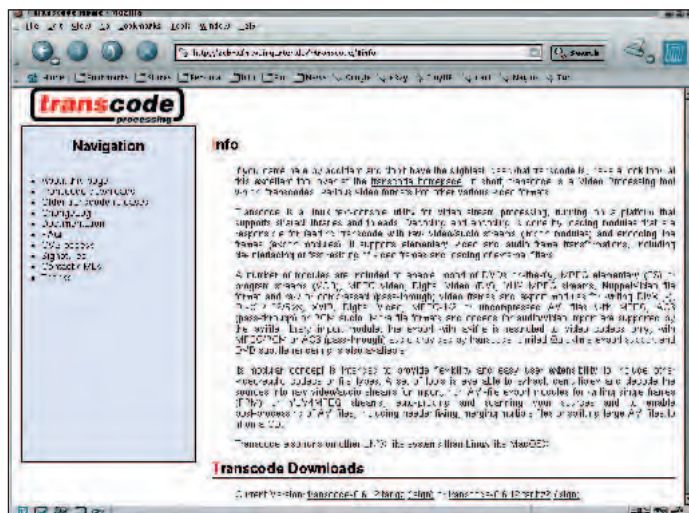
Glorious GRUB

Q I have hit a problem and now need some help – I tried to install Fedora Core 2 and Windows XP dual-boot on my laptop – a Gericom P4 with 512MB RAM and a 60GB hard drive with a Nvidia GeForce Go 5600, Smartlink 56Kbps modem and nation semiconductor 83815 NIC.

While I am having a problem with the drive geometry being changed and unable to see the Windows partition, I have a bit more serious problem in that when I install Fedora and reboot, the bootloader *GRUB* will just say “GRUB” on a black screen and will just hang there, not giving me any options. This happens anytime I try installing it as a dual-boot, and I really want Linux on my laptop as primary OS; I just have to have Windows for work applications.

Mark Mc Nicholas, via email

A The first step should be to establish why Fedora can't see the Windows filesystem. If you can sacrifice your Windows filesystem, or are easily able to back it up, it may be a good idea to start from scratch, blow away the partition table and



Transcode can create AVIs from DVDs, or create MPEG2 format video to write to a VCD, playable by many modern living-room DVD players.

recreate it exactly how you want under Linux. Leaving space for your Windows filesystem, you can then install Fedora on the clean disk and ensure that *GRUB* works happily for you.

Once Fedora works as you would wish it to, you can then do a clean Windows installation, boot off the Fedora media to restore the *GRUB* installation, then modify the *GRUB* configuration to permit you to boot the Windows system via *GRUB*.

VCD recording

Q I recently bought a Trust 782AV LCD Power Video, which is basically a glorified digital camera masquerading as a miniature video camera. The camera stores files as ASF on an SD card. According to the Trust website, the files are MPEG4 in an ASF wrapper to allow *MS Media Player* to read them.

I can access the camera's files as a mass storage device, and can successfully download these to my computer, but cannot play them. I am using Mandrake 10.0 from your coverdisc (top stuff, by the way – thanks) but cannot get the files to play using any of the videoplayer packages supplied. I have contacted Trust for advice, but the info I got was not very useful. I downloaded some codecs via the *MPlayer* site but I'm not quite sure what to do with them. I'd like to be able to edit the files and burn VCDs to play on a DVD player.

James Shaw, via email

A Certainly the best tool to use in order to convert a video to an alternative format is



Xine supports the vast majority of video formats, including AVI, MPEG and can play DVDs, including the fancy menus that preface many movies.

A QUICK REFERENCE TO: Snort

A firewall isn't the last line of defence against intruders, so here's the lowdown on IDS.

Even with a comprehensive firewall configuration, there is still the problem that we need to have at least a few basic ports open so people can actually access things on our network. Whether this is limited to *SSH* and *Apache*, or includes more complex and vulnerable processes, including *Sendmail* or *SNMP*. We can't simply block all access, however there are often exploits for CGIs or PHP scripts, or someone can attempt certain exploits on a service, such as buffer overflows on *Sendmail* or *CodeRed* against *Apache* (thinking it's *IIS*).

Monitoring traffic permitted through the firewall is the job of an Intrusion Detection System, or IDS, which continually monitors network traffic and performs packet analysis to look for

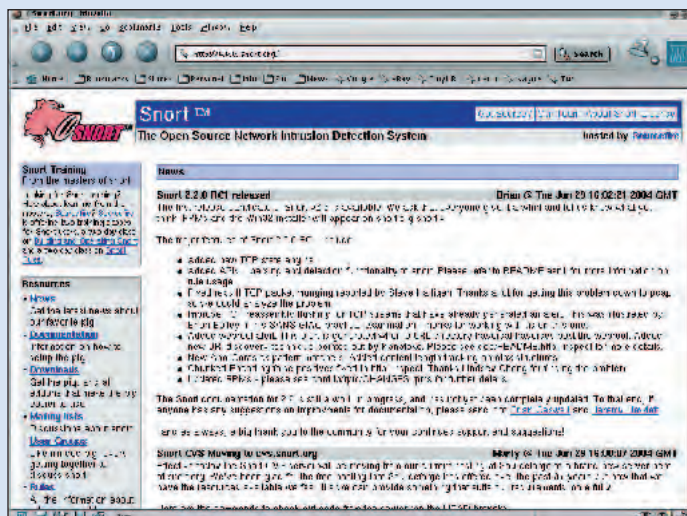
signatures from viral activity, exploits, or other unwanted happenings on our network. *Snort* is a wonderful Open Source IDS, and coupled with a Web interface such as *ACID*, or using a notification

system similar to *PigSentry* to tell us when something unwanted happens on the network.

Snort can be quite a pain to set up, though many prebuilt signatures and rules are available online, so we

don't always have to start from scratch. It's important to ensure that the rules are applicable to the environment, so if there isn't an *IIS* server, then we can just skip all of the *CodeRed* nonsense. Anyone who gets notified when there is an IDS event on a web server will really be unhappy if they get told someone tried *CodeRed* or *Nimda* against *Apache* or *Boa*! However, notification of other attacks – such as PHP exploits or file insertion exploits through CGIs – are invaluable when securing a network from the inside out.

Snort can be downloaded from www.snort.org, though most distributions will ship with *Snort*. When it comes to using it, there is quite a wealth of different sorts of documentation available on building a ready-to-roll IDS, such as the *Snort* and *ACID* HOW-TO at www.floridahoneynet.org/whites/snortacid.html.



Snort is a fantastic addition to any firewall, or network, and compliments many traditional security processes with something a little more dynamic.

transcode (www.theorie.physik.uni-goettingen.de/~ostreich/transcode/).

This program can take a variety of video sources, convert them, scale them, do pretty much whatever you want with them. It can also output the appropriate encoding for burning yourself a VCD which will play on a PAL or NTSC television.

You may be able to play ASF files using *Xine* (<http://xine.sf.net/>) which includes a variety of plugins for common formats. Out of the box, *xine* will be able to play ASF files, as well as DVDs, MPEGs, AVI and almost anything else you can think of. Being plugin-based, you can also add third-party modules supporting some more obscure formats, but *Xine* should do everything you need it to do without having to mess around too much.

Backup

Q I want to backup large amounts of data to an external DVD drive that uses a USB interface. Using *growisofs* (which uses *makeisofs*) from the command-line, I can successfully backup to the DVD-R disc. However, when I pass a

directory name as the file to be backed up, what gets copied is the directory tree underneath the passed directory name. Is there a way that the file structure can be written so that the passed directory is also written to the DVD? All subdirectories are written as they appear, just like a *cp* operation with the *-r* (recursive) switch, but the parent directory is left off.

For example:

Directory *dirname1* contains the file *file1*, and two subdirectories *subdir1* and *subdir2*. At the command-line, I enter:

```
growisofs -Z /dev/cdrom2 -J -R
dirname1
```

Everything under *dirname1* will be copied as is: *file1*, and the two subdirectories *subdir1* and *subdir2* (and their contents). *dirname1* is not copied, though.

If you then want to merge another data set – say, *dirname2* – you would then enter

```
growisofs -M /dev/cdrom2 -J -R
dirname2
```

However, my data sets are differentiated based on directory name (eg *dirname1*), so that the

same files – *file1*, *subdir1* and *subdir2* – appear in the second data set as in the first. The second backup operation will overwrite the previously written files from *dirname1*, unless the parent directory hierarchy can be preserved to the DVD.

Obviously, you could create a temp directory, then move a bunch of these directories (*dirname1*, *dirname2*, etc) into the temp, then supply the temp directory to the *growisofs* command-line. This will correctly copy the parent directories (*dirname1*, *dirname2*, etc) to the DVD, and will preserve the subdirectory structure underneath. My problem is, that the directories in my application are typically tens of megabytes each, containing approximately 1000 files each, and duplicating them (even briefly) is impractical on the hard drive. Any simple solutions?

(As an aside, part of the problem with Linux is that you never know whether there's actually something already written that you just haven't found yet, or whether you need to rewrite source

code to get the thing to do what you want it to. It's like trying to prove a null hypothesis (it can't be done), but after searching on the web for days and coming up empty handed, you come to the realisation that maybe the feature/tool doesn't exist. I would have thought that *growisofs* and *mksiofs* would behave like *CLI* move and copy operations with the *-r* switch. However, *-r* and *-R* are evidently reserved for denoting RockRidge extensions to file names.)

Tim Johnson, via email

A Very simple solution to your problem would be to use a temporary directory, as you indicate, but rather than copying everything over, one can create a link between the directory you want to backup and a directory within the temporary directory you created.

```
# mkdir /home/tmp
# ln -f /home/david
/home/tmp/david
```

You can then use *growisofs*, passing */home/tmp* as the directory with which to build the ISO image. Unlike using a symbolic link, using a 'hard' link makes */home/tmp/david* identical

FREQUENTLY ASKED QUESTIONS SYSTEM ADMINISTRATION

FAQ HOW DO I CREATE A NEW USER WITH LINUX?

The standard way to create a user is with the **useradd** command (but many distros have their own variant, such as **adduser**). Syntax is simple: **useradd [-u uid [-o]] [-g group] [-G group,...]**

[-d home] [-s shell] [-c comment] [-m [-k template]]

[-f inactive] [-e expire] [-p passwd] name

So, if we wanted to create a user called 'david' who's default group was 'users', we just need to do **useradd -g users david**

The group details are a little complex, as a user has both a default group, along with a list of other groups which they can switch to with **chgrp**. One thing to note with **useradd** is that the **-p** option inserts the supplied text directly into the **/etc/passwd** file. If we want to actually shove a working password straight in, we need to create a password hash using **crypt()**, with something like Perl or C.

FAQ CAN I GIVE THEM A DEFAULT PASSWORD?

If we're not looking at using the **-p** flag, it's easy enough to use the **passwd** program to set a password for a user. For the user 'david', we can do the following:

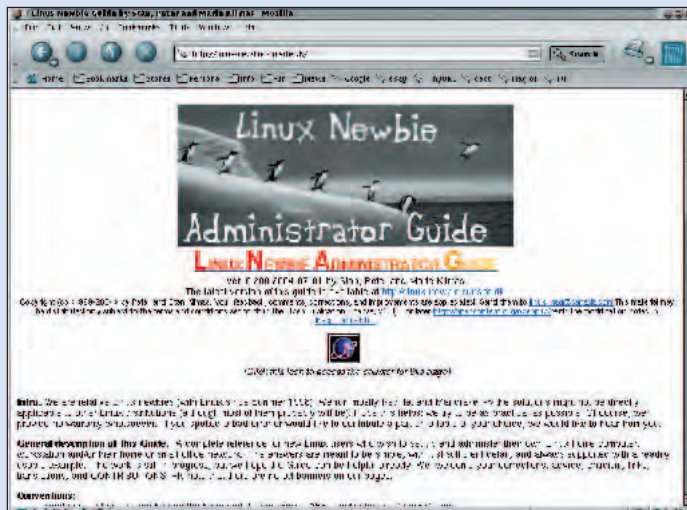
passwd david

which will ask for a new password.

Note that this should be performed as root, otherwise it will complain.

FAQ CAN I DISTRIBUTE AUTHENTICATION INFORMATION OVER A NETWORK, SO I DON'T NEED TO MAKE EVERY USER ON EACH SYSTEM?

Well, if we've got lots of free time or are being paid by the hour, we can! Unfortunately, life is not always that easy, so we can use a system known as NIS to distribute authentication information. NIS is a server/client system, so we run a NIS server which contains all of the authentication information, then on each system we



Getting started with Linux can feel like pushing string up hill, but there is so much documentation available, the answer is never far away.

run the client which permits the system to do authentication lookups via the server.

Information on NIS can be found at <http://en.tldp.org/HOWTO/NIS-HOWTO/>

FAQ CAN I STOP USERS LOGGING INTO A PARTICULAR SYSTEM, BUT STILL HAVE THEIR USER INFORMATION AVAILABLE?

When using NIS, we can modify any

to **/home/david**, in that they are both file pointers to the same inodes and directory information. You won't get any of the messy problems we see with symbolic links, where a file is referenced outside of the directory structure you want copied and the whole process gets out of hand.

When you are done, you can delete **/home/tmp/david** without any worries and create the next link to merge into your image. Note that the link must be on the same filesystem as whichever directory you want to backup, so if **/home** and **/usr** are separate filesystems, you can't link **/usr/src** to **/home/tmp/src**. However, as the link does not consume any storage space, it should not be a problem to create the link somewhere on a filesystem.

ZoneMinder

The **ZoneMinder** article in **LXF46** attracted me because it was for me a perfect article with a nice mixture of functionality and usage; so I decided to try the demo. I bought a Logitech QuickCam Express USB (a

HP Deskjet 916c is also connected through USB, both devices are using a dedicated connector at the motherboard). Running Red Hat Enterprise Server, I found that the HP printer was recognised and installed correctly. The QuickCam Express USB was not.

I checked the **ZoneMinder** support page, and found that Red Hat ES is not supported.

So, as a work-around, I installed Mandrake 9.1, but unfortunately the webcam USB is also not recognised.

The USB camera is not detected, it seems that a driver with basic functionality is all I need. I tried to find which driver was appropriate, but was overwhelmed with information, and cannot find any relationship between driver and corresponding Red Hat nor Mandrake package.

Any advice or methodology on how to start up the basic camera functionality would be appreciated! I tried also your advice from **LXF52** April 2004 - Making movies (webcam > usb storage media) error : special device does not exists

(/dev/sda1 or sdb1 or sdc1....)

Hugo Bruers, Waarloos, Belgium

The QuickCam Express USB camera is not one using a common chipset, so is not supported by the major kernel projects working on Logitech camera support. However, there is driver support using **spca5xxx** kernel modules that you can get from <http://mxhaard.free.fr/download.htm>. The Express USB camera is flagged as being 'test' only, so your mileage may vary. You may also want to check out the projects which work on support for the Logitech cameras at <http://qce-ga.sourceforge.net/> and www.smcc.demon.nl/webcam/. Note that while the qce-ga site does refer to the Quickcam Express camera, it is referring to an older model which uses a different chipset.

Firefox

I am running Mandrake 10.0, with KDE and am using Mozilla as a browser. I thought it would be nice to try Firefox as supplied on the DVD with **LXF55**. I was immediately puzzled

by the choice of three **Firefox** tar files. However, I unzipped the one with **'-PC-'** in the filename but was further puzzled when there appeared to be no **README** or **INSTALL** files etc, to provide guidance about installing.

After a bit of research, I found that there is probably a bit of script in an **.sh** file which needed to be run and there is indeed a file called **run-mozilla.sh**. However, my research failed to turn up any explanation of how to 'run' the script. In the end, I just entered the path and name of the file in a terminal (as root) and pressed enter. The response I got was "Cannot Execute," sadly, with no further explanation.

I then tried the other tar file (not the one with 'source' in the name) with the same result. Is it me, or is it the software?

Peter McPherson, via email

The latest stable versions of **Firefox** and **Thunderbird** are both available on this month's **LXF** coverdiscs for both DVD and CD readers. **Firefox** is distributed as a

/etc/passwd information, such as the shell. To only permit users from the **admin** netgroup to shell into the host, we can do this:

```
+@admin::0:0::
+::0:0::/bin/false
```

FAQ HOW DO I SET STORAGE QUOTAS FOR EACH USER ON /HOME?

Most file systems – including ext3 and ReiserFS – have quota support, which requires us to compile quota support into the main kernel. By mounting the /home file system with the **usrquota** option, we can give each individual user their own storage quota, which is based on both bytes used as well as inodes consumed.

The command **repquota** displays the current quota information for all users on the system, as shown in the table on the right:

As we can see, there is both a **soft quota** and a **hard quota**. A soft quota can be exceeded, but the user is emailed every day about being over the limit. A hard quota is fixed, and the user is unable to write any more data than this limit.

For each user, we can run **edquota <username>** and modify the file as appropriate. If lots of users are going to have the same quota information, we can start by creating a prototype user quota, then replicate it with:

```
edquota -p <protouser> david
```

FAQ CAN I PLACE A LIMIT ON THE NUMBER OF PROCESSES OR MEMORY THAT MY USERS CAN CONSUME?

ulimit can be used to limit the number of processes a user can have running, and the total amount of RAM to be used by those processes. An example

usage of **ulimit** would be:

```
ulimit -u 35
ulimit -s 2048
ulimit -c 0
ulimit -m 25600
```

The first option limits the number of processes to **35**. It's best not to set this too low, otherwise the user won't be able to actually do anything. **-s** limits the stack size to 2MB and **-c** sets the maximum core size, which is the dump following a segfault, to zero; meaning that no core file will be generated. Finally, in the last line here we limit the amount of RAM used by the user to **25600KB**.

FAQ WHAT ABOUT CPU USAGE? CAN I LIMIT THE AMOUNT OF CPU TIME A USER OR PROCESS CAN CONSUME?

nice can be used to limit process priority at runtime, and **renice** can be used to modify a running process priority. However, this is all open to abuse by users, so an alternative method should be considered.

Instead, we can apply the *fair scheduler* patch which reduces process priority when a particular user is using too much CPU, effectively creating a dynamic **renice** system for the processes. This patch can be found at www.surriel.com/patches/

REPQUOTA

```
maeve:~# repquota -a
```

```
*** Report for user quotas on device /dev/maeve/home
```

```
Block grace time: 7days; Inode grace time: 7days
```

User	Block limits				File limits			
	used	soft	hard	grace	used	soft	hard	grace
root	13	0	0		4	0	0	
debian	29	0	0		9	0	0	
david	5881991	0	0		212568	0	0	



Firefox is a great web browser for those who don't want to deal with the 'everything including kitchen sink' approach of the whole **Mozilla** package.

complete app that you can execute with the **run-mozilla.sh** script. You should be able to **cd** into the untarred **Firefox** directory and do **./firefox** to get it all up and running. As a more permanent solution, one can move the 'firefox' directory to /usr/local/firefox, create a symlink from /usr/local/firefox/firefox to

/usr/local/bin/firefox and then run it from the command-line with **firefox**.

VCR backup?!

Q I have been waiting for some time for the Audigy2 soundcard to be implemented fully under Linux (ALSA), and with

kernel 2.6 and the new round of distros this was delivered. But when I installed Mandrake 10.0, I was having similar problems to David Kendall (*Answers, LXF55*), with the sound support not starting correctly.

What I found after I sussed that the driver that is supposed to run my card, an Audigy, didn't and it was **snd_emu10k** that actually works on my system, is that the sound applet **KMix** seems to grab the card before ALSA has initialised. So what I found seems to work is: quit **KMix** first, then go to 'System' > 'Configuration' > 'Configure your Computer' then choose the 'System' icon; then 'Services', then from 'Services' start **ALSA**. Something similar also works under **GNOME**.

Something else occurred to me when I was messing around with my new digital camcorder and your feature on digital video in *LXF52*. You state that raw video takes up a great deal of space, so camcorder tapes must be quite high capacity; and with the IEEE1394 being very fast and able to control the camcorder's tape transport fully,

there must be scope for using these camcorders as back-up facilities. What with tape back-up systems being quite expensive and backing up your data regularly being important this might be a cheap way of utilising a piece of equipment that people already own. I would imagine that most of the routines exist; it would just be a matter of re-rigging them so that different file type could be written to the camcorder or something?

Andy Mackay, Glasgow

A Your system should start **ALSA** through an **init.d** script, which loads the appropriate kernel modules and makes everything work happily, prior to getting close to starting up KDE. If you installed **ALSA** from the source code, you will have to locate an appropriate **init.d** script for your distro.

There is the beginnings of a project to do exactly what you suggest using a DV Camera called *backfire* (www.ajwm.net/backfire/). It seems to have been discontinued for a while; however, someone from the IEEE1394 may have reincarnated the application in recent years.

ANSWERS



Linux in training

Q I work at a computer training college where we have various courses that rely on different programs and applications for the various different courses, but most of these courses are done on the same machines. So to save time, with Windows I would just install the

relevant programs I wanted for a course: for example, if it was web design, I would install *Adobe Photoshop, Dreamweaver and Flash*, and then remove certain registry keys from the registry that contained the hardware. Then I would use *Norton Ghost* to make an image of the hard drive and name it 'web design.'

That meant when it came to that class, instead of reloading Windows and reinstalling the apps, I could just load the image, and when Windows booted, it would re-install the hardware as if it was the first time. Saving you money, saving you time and making me look smart. But now we have decided to do all the Microsoft courses, and create Linux equivalents. My question is, is there a way of performing the same tasks and getting the same result for Linux? *Warren, via email*

A You'll be happy to know that this should be very straightforward under Linux. You can still use Norton's *Ghost* to copy Linux drives as it support the

ext2 and ext3 filesystems. There are other ways to mirror these drives from within Linux but, if you're not a purist and already own *Ghost*, you may as well use it. If you are using DHCP to assign IP addresses, you should be able to move the drive on your system to any similar system with minimal changes. Pretty much all of the out-the-box distributions build kernels, which support most common hardware. Some also use a daemon to monitor changes to hardware (*kudzu* in Red Hat) – this will detect any hardware changes since the last bootup, and prompt you for how to handle these changes. If your distribution of choice does not monitor for hardware changes, you should still not have too much trouble,

★ Star Question – AV140 winner!

This issue's lucky is **Damian** – your prize will be with you shortly!

Q Love the magazine – I've got all the issues (and I've been a subscriber since quite early on). I'd like to take this opportunity to ask a few questions: **1** I work in the IT department of our local hospital – in fact, I *am* the department! I've written a helpdesk in *MySQL/PHP* running on Linux Mandrake 9.2 – the helpdesk started mainly as an exercise in learning PHP. I've got an iPAQ HP3970 (running Familiar) and would like to be able to use that when I'm out and about. Is it possible to run PHP scripts on an iPAQ? I don't have wireless connectivity, which is why I need to run the scripts on the iPAQ. Is it possible to run a webserver on a handheld for this purpose?

2 Can you recommend any good Free/Open source asset management agents? I'd like to be able to use the network to collect as much information as I can about all our PCs (over 900 networked machines)?

3 I enjoyed your Open Source Java feature in *LXF54's Linux Pro*. It raised a question about GPL: leaving aside the arguments why they wouldn't; what, for example, would happen if Microsoft *did* take some GPL code, modify it and refuse to give the modifications back to the community? I guess

I'm asking what would the community do in such a situation? What could they do?

4 Is it possible to physically trace a device on the network using just its IP address? I've been told there's a device on our network that is broadcasting lots of packets and is causing problems for our Radiology department.

Unfortunately, the people who told me about the device could offer no further help. Our PCs use an IP range 141.97 but this device is 10.0.0.12. It doesn't respond to pings nor can I *telnet/ftp* to it. *Traceroute* doesn't find it; *Nmap* seems to be able to see the host, but offers no information. It's just occurred to me to try *ethereal* to see if there's anything in the packets, which might give the game away, but I was wondering if you could offer any suggestions. (I'm not asking you for detailed instructions, any pointers you could give me would be great!)

Damian, via email

A Thanks for the long list of questions Damian. We're not sure we'll be able to answer every one, but we'll do our best to at least point you in the right direction.

1 To be honest, we have very little experience with Familiar Linux – we don't even have an iPAQ to hand in order to try it – but we'll feature

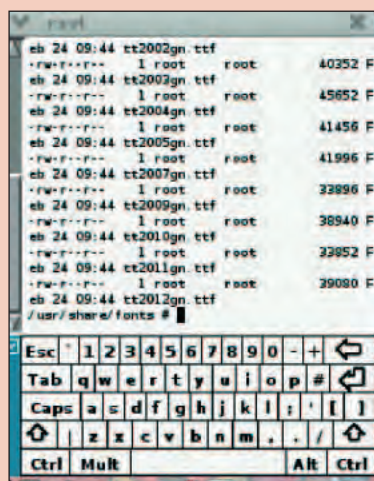
something about this in a future issue if enough of our readers request it. However, we can confirm that there is an ARM architecture .ipk available for PHP 4. Browse on over to <http://ipkgfind.handhelds.org/> and search for 'php'. From taking a look at the Familiar site, we can see that it's already a pretty large project and we're confident there will be a working Apache+PHP implementation

2 We've managed to find a piece of Open Source software, called *nextB*, which seems to be able to do what you want it to. It may be worth checking out the project's site at <http://sourceforge.net/projects/nextb>

3 It would be entirely up the author of a piece of GPLed software to pursue any license infringements on their software. Unfortunately, since most Open Source developers do what they do for the love of it (not for the money), they don't have the resources to take a company such as Microsoft to court. Many legal experts believe that the GPL is enforceable in a court of law (and just as many would say it is *not*, too), but there has been no landmark court case yet to test this. Just recently, Fyodor – the author of the Open Source *nmap* program – revoked SCO's right to

distribute his software. To this day, SCO is distributing it with its OS. There are non-profit organisations, such as the Electronic Frontier Foundation (www.eff.org), who might be able provide assistance (both financial and professional) in such cases.

4 The only completely reliable way to trace this device physically is to gather information from the network device it is attached to. Any decent manageable network switch should allow you to determine the port a MAC address is associated with. It's then a matter of tracing the cable to its physical location. You *did* label those cables, right?



The Linux command-line on an iPAQ! For more Linux PDA thrills, see page 18.

WIN an Archos AV140 Video Recorder WITH RACKSPACE MANAGED HOSTING

www.rackspace.co.uk

Every month, the best question related to Systems Administration that a LXF reader sends in wins a prize: this month, you have the chance to win the awesome Archos AV140 Video Recorder.

The AV140 combines a modular MP4 video player and recorder, MP3 music player and recorder, digital camera and camcorder, digital photo wallet and data storage in a compact device that fits in the palm of your hand and is compatible with Linux. Based on a 40GB hard disk that is completely compatible with Linux, Windows and Mac OS, all you have to do is connect it to your computer and it will be mounted as an additional hard disk – on which you can store any type of data. USB 2.0 auto-

connection ensures fast data transfer to and from your PC.

You can now enjoy 160 hours of near-DVD quality MP4 video with MP3 sound wherever you are. You can record and playback to your TV, VCR or camcorder directly in MP4 format. Playback is an AVI format file with MPEG4-SP VHS quality video and MP3 stereo sound, with the ability to record up to 2,000 hours of voice audio. See www.archos.com/products/prw_500432_specs.html.

Email all your sysadmin questions and problems to our expert system administrator Hans Huberland:

sysadminqa@rackspace.co.uk

Every published query on these system admin pages wins a Rackspace T-shirt!

WIN!



the most common hardware to cause a problem would be the network card and to resolve this change to `/etc/modules.conf` (or similar depending on your distro). The file should resemble the following:

```
alias eth0 3c59x
```

This indicates that you have a 3Com 3c59x series adapter. If the new machine has an Intel NIC for example you could just change this to

```
alias eth0 e100
```

Username ululation

Q How do I get Linux to accept usernames like "sherman.lilly"? I can't get my system to accept usernames that have a "." in them.

Sherman, via email

A The period "." character is a special character in Linux, and therefore should not be used in a username. It is possible to force this but you will need to add the user manually to `/etc/passwd` and `/etc/shadow`. Alternatively, you could add the username without the . and then modify `/etc/passwd` and `/etc/shadow` afterwards to add the . in. We've not experienced any problem doing this, but some programs could have problems because they are not expecting this character.

If you only want to add the period for the email address, then we would

recommend using an alias to do this and pass the mail to a technically correct username.

Sendmail function

Q In *qmail*, you can put an address in the *badrcptto* file and *qmail* will reject the connection based on this address without receiving the message. This

is especially useful when you have an address that no longer exists and still gets spammed.

Is there similar functionality available in *Sendmail*?

A User, From the Rackspace Forums

A You can add an entry to `/etc/mail/access` as follows:

```
user@domain.com  
REJECT
```

This will just issue a generic error. If this is intended for a user who has left it may be better to use a more descriptive error:

```
user@domain.com ERROR:"550  
User no longer exists at this domain"
```

You'll need to run **makemap** to rebuild the database file from the flat text file when you're done:

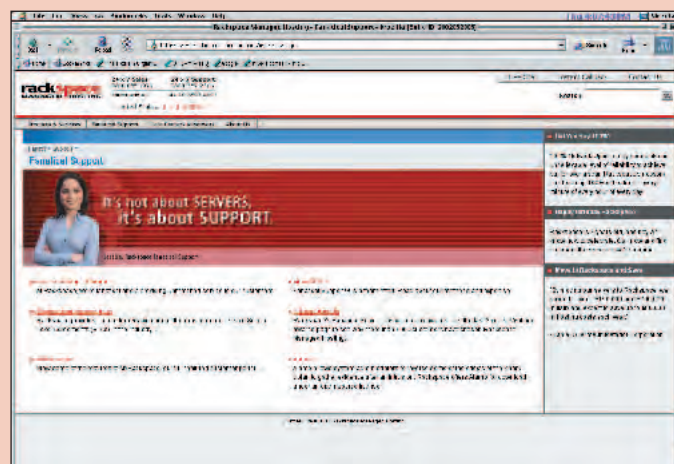
```
makemap hash access < access
```

ALL ABOUT RACKSPACE

When it comes to managed hosting Rackspace is unique. We are the only hosting company to guarantee 100 per cent network uptime and we will even pledge to replace faulty hardware (let's face it, it does happen) within an hour.

As you would expect from Red Hat's only Advanced Hosting partner you will receive Fanatical Support™ and instant emergency response from fully qualified level three technicians, available 24-hours a day, 365 days a year. No

answering machines, no diversions, no silly on-hold messages. All our customer platforms are housed in state-of-the-art, secure data centres and can be fully customised to meet any requirement. A dedicated account manager is allocated to you at the outset and will remain on your account. All Rackspace employees are committed to ensuring customer satisfaction, and will not rest until a client problem is resolved. Since Rackspace was established in 1998, it has specialised in managed hosting – nothing else. We're not an ISP, we don't sell domain names, and we won't offer website design – all we do is managed hosting. As a result, we do it very well and 97 per cent of our customers would happily recommend us. For information on how we can improve your web hosting, please contact us at www.rackspace.co.uk and ask any questions you may have about the ultimate managed hosting service.



« Device drivers

Q I am writing a driver with considerable help from the book *Linux Device Drivers* (2nd Edition) by Alessandro Rubini and Jonathan Corbet (O'Reilly, ISBN 0-59600-008-1). When I define `__KERNEL__`, the line `#include <linux/kernel.h>` produces an error from `gcc 2.95.3`. (See below.) This error does not occur when `__KERNEL__` is not defined.

I am using kernel 2.4.19 which I compiled without any problems. I see that `stdargs.h` is included when the error occurs – extension of `gcc` to functions with variable number of arguments? This is the only include file in the `__KERNEL__` section of `kernel.h` not stored with the kernel headers.

The puzzle is:

- Compiling the kernel works, `gcc` accepts `kernel.h`
- When compiling a module, `gcc` fails `kernel.h`
- `stdargs` exists in the `gcc` library
- `stdarg.h` must be found as it is not reported as absent
- `gcc 2.95.3` is the recommended compiler for kernel 2.4.19

Compiling the following is all that is necessary to produce the error and to show that it is not a problem with the coding of the module:

```
test.c:-
#ifdef __GNU_SOURCE
#define __GNU_SOURCE
#endif
#ifdef MODULE
#define MODULE
#endif
#include <linux/kernel.h>
int main(){return 0;}
```

Result:

```
root@localhost]# gcc -Wall test.c
root@localhost]#
```


This creates `a.out` without any problem.

```
root@localhost]# gcc -Wall -D__KERNEL__ test.c
```

In file included from `test.c:7`:

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 Linux discussion going on!



A helpful reference to have around – visit <http://safari.oreilly.com/> to see the wealth of information you can access online for a fixed fee.

```
/usr/include/linux/kernel.h:74: parse
error before `size_t`
```

```
/usr/include/linux/kernel.h:76: parse
error before `size_t`
```

```
root@localhost]#
```

Compilation has failed!

I have seen a number of bug reports referring to the apparent absence of `stdargs.h`, but none which refer to the above problem. I have tried contacting the maintainers of `kernel.h` without success, understandably because it works with the kernel!

Peter Antonelli, via email

A It looks like your headers are mixed up for some reason. We would suggest verifying that `/usr/include/linux` is a symlink to your `/usr/src/linux/include` directory, so it is referring to the appropriate include path. You can also do this by passing a `-I` flag to `gcc` to include a specific directory, like this example:

```
# gcc -Wall -D __KERNEL__ -
I/usr/src/linux/include test.c
```

Webcam & Audio

Q I have friends that I keep in touch with by chatting online. Since I am still learning about Linux, I use *MSN Messenger* to do this, and a lot of the time we end up using the audio and webcam capabilities to talk to each other. I have found *Kopete*, *GAIM* and others for Linux, but none seem to support audio and webcam features (unless I just wasn't paying attention). I've found *GNOME Meeting*, but do my friends have to be using *GNOME Meeting* as well for us to communicate, or can it work with other apps? Is there

something else out there that I might try for this?

Scott Johnson, via email

A *GNOME Meeting* is a standard H.323 client, so it will work quite happily with *Netmeeting* without a great deal of fuss. There is also a Linux *Netmeeting* HOW-TO, which you can find at www.freesoft.org/software/NetMeeting/, which uses *ohphone* as example software. Certainly, out of the Linux options for a H.323 client, *GNOME Meeting* is the most robust and well maintained. It may be worth trying out *GNOME Meeting* and seeing if it is missing anything specific which you require. From there, you can look at development versions of the software, or make suggestions to the developers as to how they can improve the package, and you could even help with writing some documentation!

Virtual virus


Q During the last three years, the Linux kernel has been a rock. My pc runs 24/7 also acting as a NIS, file and mail server for the kids' machines.

I know that users' choice of distributions is a very personal thing, but I've really grown to love my Slackware/Dropline GNOME setup. The only thing that worries me about this is that I have become reliant on downloading binaries to keep my system up-to-date, because compiling GNOME from sources is a pain in the rear. So how safe is this practice and how would I know if a hacker had planted something nasty at dropline.net which I then install on my system as root?

I figured that the best I could do would be to scan for viruses after each upgrade so I have installed *ClamAV*. So far, I haven't found any viruses lurking around, apart from *Exploit.IFrame.Gen* inside a *VMware* Windows 98 virtual disk file. Do viruses in this sort of file pose a risk, or can I ignore them?

Steve Roper, via email

A The best tool for verifying a Linux system is the popular *chkrootkit*, which performs a variety of process and binary validations in order to see if your system has been compromised; and if someone has installed a rootkit or kernel module to maintain access to the system following you solving the problem which caused the security breach in the first place. See www.chkrootkit.org/ for more info.



chkrootkit should be a standard installation on any Linux system, and can be easily run from a cronjob nightly to check for security issues.

When installing packages or binaries, one should verify MD5 information from the source, such as a post on a mailing list or from the website, against what was downloaded from the mirror. This will verify that what you downloaded is exactly what the person uploaded to the website when they built the package or tarred up the source.

Within a *VMware* image, the virus can't cause any problems to your Linux system, but can certainly cause the same problems within *VMware* as it would on a regular Windows box. *VMware* is no excuse not to install up-to-date virus-scanning tools and not keep the Windows Updates current with the latest patches and service packs from Microsoft. **LXF**

Submission advice

We are happy to answer all sorts of Linux-related questions. If we don't know the answer, we'll find out for you! But in order to give you the best service, it helps a lot if you read the following submission advice.

- Please be sure to include any relevant details of your system. "I can't get Z to work" doesn't really mean anything to us if we don't know things like what version of Z 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 to answer most questions. If we don't answer yours specifically, you'll probably find we've answered one just like it. We regret that we can't 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

missed one?

LINUX FORMAT BACK ISSUES

Every month *Linux Format* brings you the best tutorials, the essential reviews and the latest news. But if you've missed out on a must-read feature or a vital program from our expertly compiled CDs and DVDs, order your back issue NOW! And remember, you need never miss an issue of your favourite Linux mag, if you subscribe to *Linux Format* (see overleaf for more details).




August 2004

Product code:
LXFB0056(cd)
LXFD0056(dvd)

DVD HIGHLIGHTS:
PearPC, Wine, Tagged
Message Delivery Agent,
Mount ISO Image

CDs HIGHLIGHTS:
Fedora Core 2, GCC
3.4.0, Astaro Security
Linux, Blender 2.33a,
Nvu, Chatzilla, BitchX,
KVIrc, TalkSoup, Irssi,
XChat, RSS software,
SuperTux (Mario clone)

MAGAZINE FEATURING:
Fedora Core 2, Red Hat, Win
a £2K Dell server, Patent
Insanity, Inkspace, Postal 2,
GCC 3.4.0, CrossOver Office,
Recode tutorial, Beginners'
guide to FTP, LUG Infopoints



July 2004

Product code:
LXFB0055(cd)
LXFD0055(dvd)

DVD HIGHLIGHTS:
OpenOffice.org 111,
Debts, KDeTV, XROG-
X11, XFree86, DSPAM,
Abcde, GnuCash, XRoar

CDs HIGHLIGHTS:
Dyne:bolic GNU-Linux,
Cooperative Linux,
KDevelop, Leylines,
Mozilla 1.7, W3M, Firefox,
Konqueror, Links, Lynx,
Epiphany, Galeon

MAGAZINE FEATURING:
Love code: a beginners' guide,
Neverwinter Nights, Maya 6,
LDAP basics, Web browsers,
the O(1) Scheduler, GNOME
2.6, Gigabyte SR147, KDevelop,
hdparm, Kopete, GIMP coding



June 2004

Product code:
LXFB0055(cd)
LXFD0055(dvd)

DVD HIGHLIGHTS:
Live Distro Bonanza!
Seven ready to burn –
including MandrakeMove,
plus Knoppix boots
directly from the DVD

CDs HIGHLIGHTS:
Knoppix, Damn Small
Linux, Slax, Storix
Personal, RPM Linux,
GNOME 2.6, Security
applications, The GIMP 2

MAGAZINE FEATURING:
GET PROTECTED – combating
your PC problems, Creative
Commons, VMware GSX,
Audacity, Java 2, TuxCards,
Pedora Core 2, SUSE Pro 9.1,
KDevelop tutorial: Part one



May 2004

Product code:
LXFB0053(cd)
LXFD0053(dvd)

DVD HIGHLIGHTS:
Komics, K3b, Unison,
Unreal Tournament
2004, Subversion

CDs HIGHLIGHTS:
Mandrake 10.0
(including Kernel 2.6.3,
KDE 3.2, Mozilla 1.6,
OpenOffice.org 1.1.0,
GIMP 2, Samba 3.0,
MySQL4 etc) Kino,
Subversion, NET-SNMP

MAGAZINE FEATURING:
Build the ULTIMATE Linux PC,
Mandrake 10.0, C/C++ IDEs,
Tim O'Reilly exclusive interview,
UT2004, SQL databases,
Mono, Evolution, Subversion,
GIMP 2 preview, KDE 3.2



April 2004

Product code:
LXFB0052(cd)
LXFD0052(dvd)

DVD HIGHLIGHTS:
8 Distros, Gcompris,
MyPasswordSafe,
SystemRescueCD

CDs HIGHLIGHTS:
Tons of video apps,
KDE 3.2, ClarkConnect,
SMEServer, LMSensors,
hdparm, LiarLiar,
QTVision, SpamAssassin,
AtomicTanks, ifplugd,
Roundup51

MAGAZINE FEATURING:
Making movies, Opteron or
Itanium? Compression utilities,
HDL, SDL game programming
tutorial, Perl Template Toolkit,
KOrganizer, Eric Raymond,
SNMP, Arkeia 5.2 review



March 2004

Product code:
LXFB0051(cd)
LXFD0051(dvd)

DVD HIGHLIGHTS:
Java Desktop System,
8 mini distros, Scrubber,
Python 2.3.3, Bacula

CDs HIGHLIGHTS:
MandrakeMove,
Seapine Surround SCM
2.1, Simple CDR-X,
MythTV, Six mini distros,
Konserve, Tuxpaint,
Kernels 2.4.24 & 2.6.1,
Hardware applications

MAGAZINE FEATURING:
Get your hardware working,
Tadpole laptop, Xandros OS2,
Firewalls Roundup, PixiePlus,
Armari 4-way Opteron, Xinit
SPS440, Wine Rack, Opteron
vs Itanium2, MandrakeMove

Other back issues also available, please call for further information, giving the issue number and/or date required

To order, call the back issues hotline on **0870 4448645**

or **+44 1458 271178** for overseas orders

Remember to quote the issue code shown above and
have your credit card/switch card details ready.

Prices (inc post and packing): CD – UK£5.99,
US & Europe £6.99, Rest of World £7.99
DVD – UK £6.99, US & Europe £7.99, Rest of World £8.99.

LINUX

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

PRIORITY ORDER FORM:

☐ I would like to subscribe/renew to *Linux Format* (please delete as appropriate) and get my free copy of *The Complete Linux Handbook 2*

Select CD or DVD option:

	UK	Europe/N.America	Rest of World
<input type="checkbox"/> CD	£54.90	£78.90	£91.90
<input type="checkbox"/> DVD	£59.90	£83.90	£96.90

YOUR DETAILS

Title _____ Initials _____ Surname _____

Address _____

Postcode _____ Country _____

Email address _____ ☐ business ☐ home

Telephone number _____ ☐ business ☐ home

Mobile _____

Your subscription will start with the next available issue

We will use the contact details supplied to communicate with you about your subscription. Indicate here if you're happy for *Linux Format*, Future Publishing and its sister companies to contact you by: ☐ email ☐ mobile. If you're also happy for us to pass your details on to carefully selected companies indicate here: ☐ email ☐ mobile

PAYMENT DETAILS [please tick]

I understand that I will receive 13 issues during the next 12 months.

☐ I enclose a cheque for _____ made payable to Future Publishing Ltd

I wish to pay by

☐ Visa ☐ MasterCard ☐ American Express

☐ Switch: issue no./valid date _____

Card number

Expiry date _____

Signature _____ Date _____

☐ Direct Debit (UK only – double CD Edition) – I wish to pay £13.73 every 3 months

or ☐ Direct Debit (UK only – DVD Edition) – I wish to pay £14.98 every 3 months

Please complete the form below.

Details of Direct Debit Guarantee are available on request.

 Media with passion Future Publishing Ltd, Cary Court, Somerton, Somerset, TA11 6BR	Instruction to your Bank or Building Society to pay Direct Debits Originator's Identification Number 768195	
	Please send this form to address shown	

1. Name and full postal address of your Bank or Building Society branch To: The Manager _____ Bank/Building Society Address _____ Postcode _____	
2. Name(s) of account holder(s) _____	
3. Branch sort code (from the top right-hand corner of your cheque)	<input type="text"/> <input type="text"/> <input type="text"/> <input type="text"/> <input type="text"/> <input type="text"/>
4. Bank or Building Society account number	<input type="text"/> <input type="text"/> <input type="text"/> <input type="text"/> <input type="text"/> <input type="text"/> <input type="text"/> <input type="text"/> <input type="text"/> <input type="text"/>
5. Instruction to your Bank or Building Society Please pay Future Publishing Direct Debits from the account detailed on this instruction subject to the safeguards assured by the Direct Debit Guarantee. I understand that this instruction may remain with Future Publishing and if so, details will be passed electronically to my bank or building society.	
Signature(s) _____ Date _____	
Ref No (Office use only) _____	

You may photocopy this blank form and then complete it. You cannot fax or email Direct Debit forms.

UK READERS: *Linux Format*, Subscriptions, Future Publishing Ltd, Freepost BS4900, Somerton, Somerset, TA11 6BR

OVERSEAS READERS: *Linux Format*, Subscriptions, Future Publishing Ltd, Cary Court, Somerton, Somerset, TA11 6TB, United Kingdom

Please indicate if you do NOT wish to receive relevant information about special offers or products/services from *Linux Format*, Future Publishing or any of its sister companies by ☐ post or ☐ telephone. Please indicate if you do NOT wish us to pass your details on to other carefully selected companies to enable them to contact you about their products/services by ☐ post or ☐ telephone.

CLOSING DATE 01.09.04

OFFER CODE LXF57



SAVE OVER £35!†

LINUX FORMAT BRINGS YOU:

- >> Essential reading for all Linux users
- >> 20-page *Linux Pro* supplement with every issue
- >> CD or DVD packed with Linux software
- >> Latest news, latest reviews and the best Linux tutorials around

3 EASY WAYS TO SUBSCRIBE:

- >> Call 0870 444 8645
- >> Post Future Publishing Ltd, Freepost, TA11 6BR
- >> Fax your completed form to: 01458 271146
- >> Online at www.linuxformat.co.uk

Offer ends 01/09/2004

*Gift stock is limited. In the event of this gift not being available, Future Publishing reserves the right to offer a substitute gift of a similar value. Please allow 28 days for delivery of your gift.

† Saving based on 13 issues of the DVD edition, bought on the UK newsstand, plus *Linux Handbook 2* worth £14.99




FREE GIFT!

Subscribe today and get *The Complete Linux Handbook 2* absolutely free*

Subscribe to 13 issues of your favourite Linux magazine and get *The Complete Linux Handbook 2*, worth £14.99, absolutely free!*

KICK-START YOUR LINUX EXPERIENCE!

- Get started with Linux
- Replace costly applications
- Edit your images
- Set up a network
- Learn from the experts
- Unbeatable coverage



**RRP
£14.99!**

ON YOUR DVD...

- Linux Mandrake 9.2 – complete and updated OS
- OpenOffice.org 1.1.0 – full-featured office suite
- Hottest 50 tools ever, as chosen by *Linux Format*
- CD burners, web servers and astronomy tools

ORDER HOTLINE ☎ 0870 444 8645 • Please quote LXF57 when placing your order
OVERSEAS CUSTOMERS CALL: +44 870 444 8645

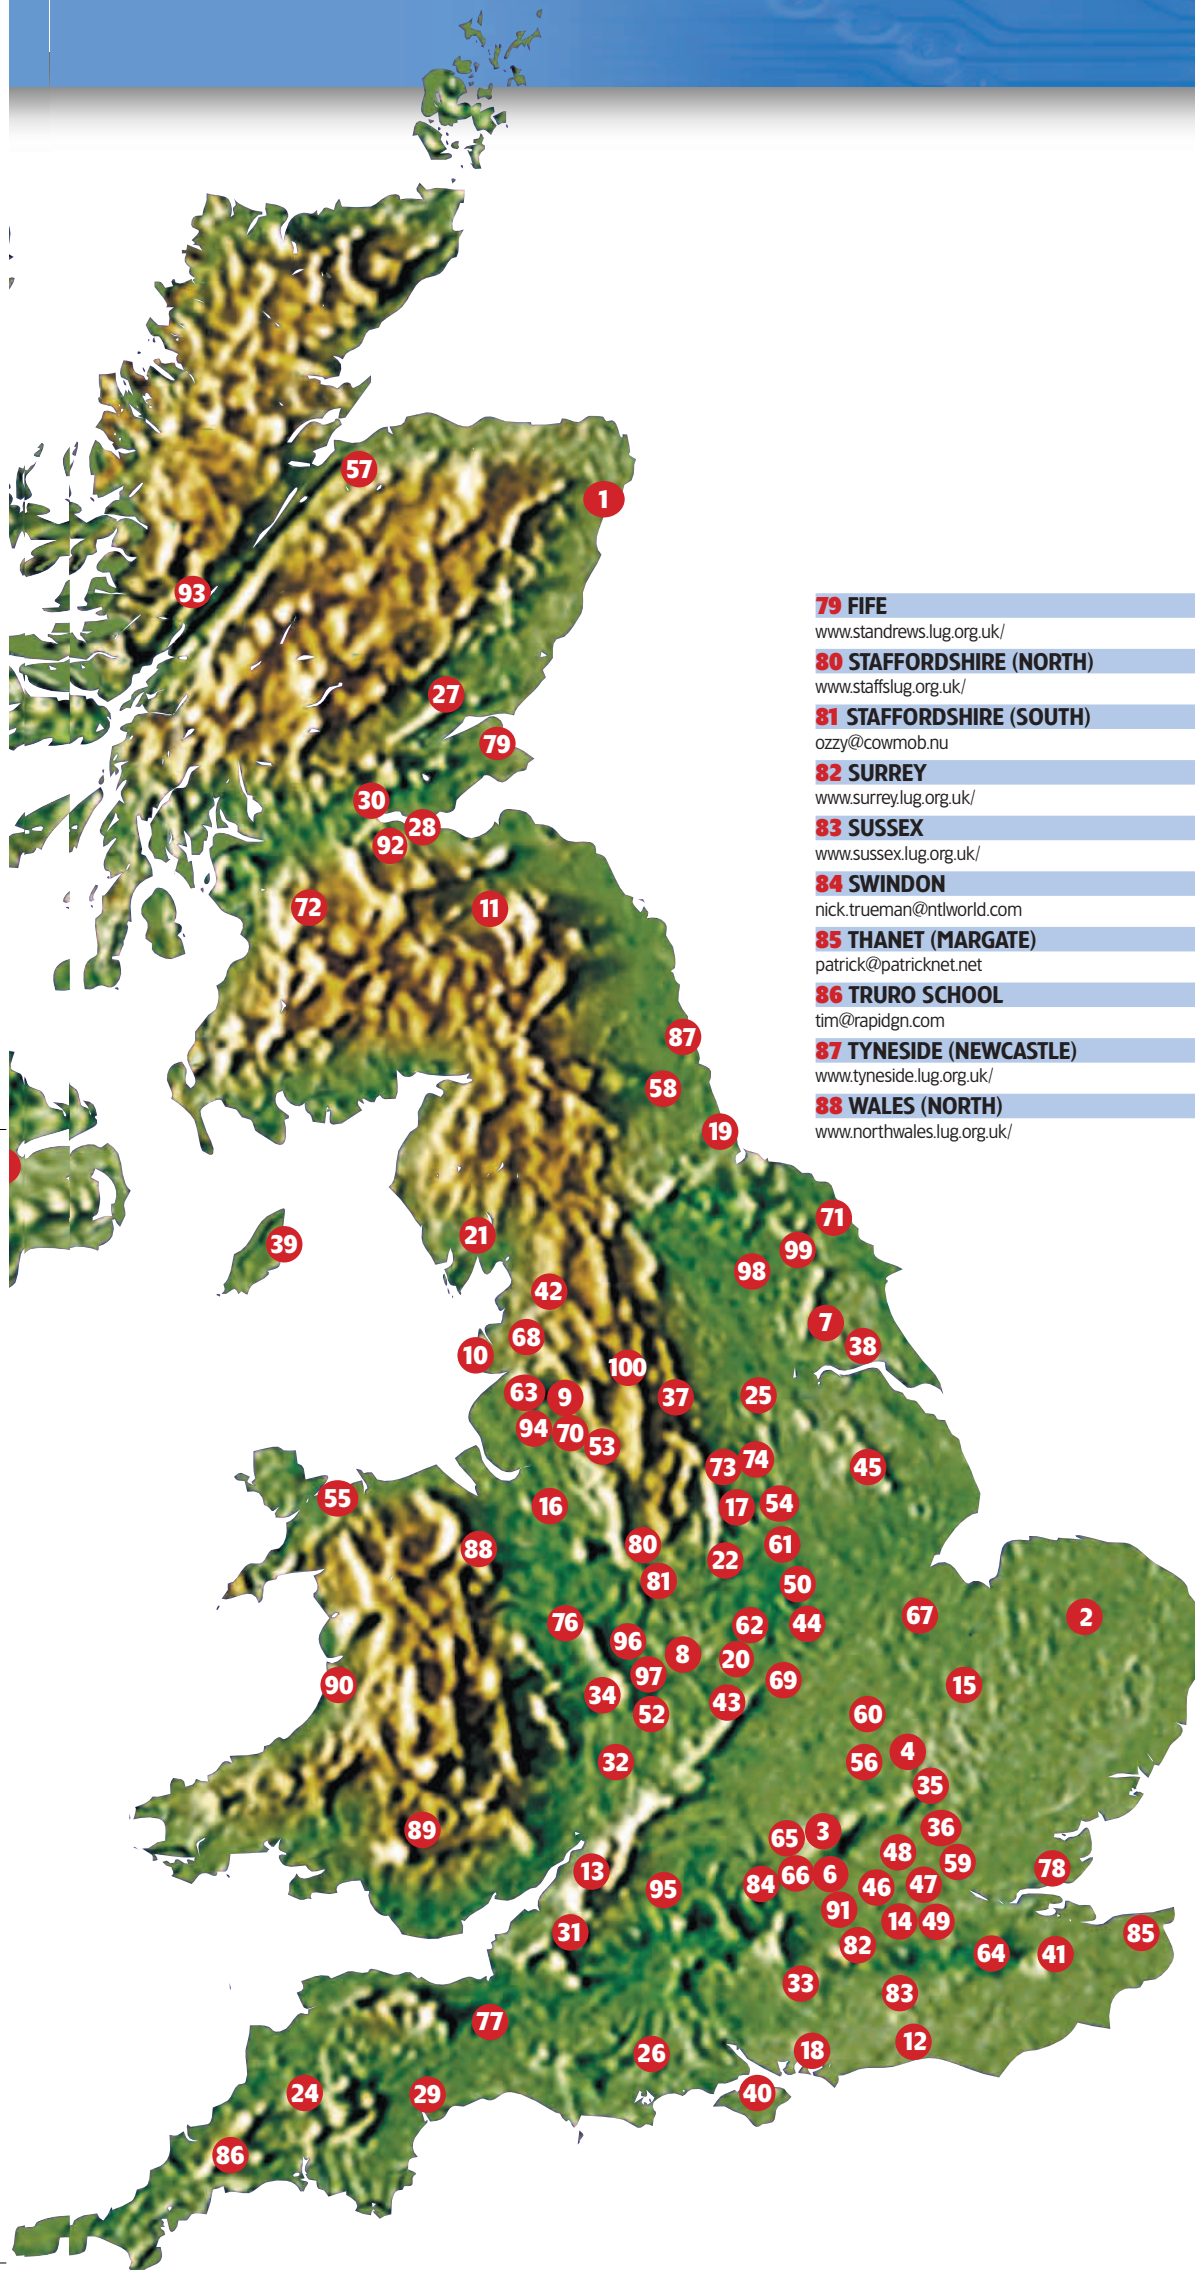
User Groups

LUGs worldwide are full of members keen to help with your problems, discuss ideas, and generally chat about all things Linux. You can find lots more information online at www.lug.org.uk

1 ABERDEEN www.aberdeen.lug.org.uk/	27 DUNDEE & TAYSIDE www.dundee.lug.org.uk/	53 MANCHESTER www.manlug.mcc.ac.uk/
2 ANGLIAN LUG www.alug.org.uk/	28 EDINBURGH www.edlug.org.uk/	54 MANSFIELD www.mansfield.lug.org.uk/
3 AYLESBURY drbond@educational-computing.co.uk/	29 EXETER (UNIVERSITY) nick@urgusabic.net/	55 MENAI kevin@dotmon.com
4 BEDFORDSHIRE www.beds.lug.org.uk/	30 FALKIRK www.falkirk.lug.org.uk/	56 MILTON KEYNES www.mk.lug.org.uk/
5 BELFAST & N. IRELAND bgalbraith@gmx.net/	31 GLASTONBURY www.glastonbury.lug.org.uk/	57 MORAY Stu@forres.co.uk
6 BERKSHIRE & THAMES VALLEY www.sclug.org.uk/	32 GLOUCESTERSHIRE & COTSWOLDS www.gloucs.lug.org.uk/	58 NORTH EAST ENGLAND (DURHAM) www.ne.lug.org.uk/
7 BEVERLEY www.beverley.lug.org.uk/	33 HAMPSHIRE www.hantslug.org.uk/	59 NORTH LONDON jason@voyagercomputers.co.uk
8 BIRMINGHAM (SOUTH) www.sb.lug.org.uk/	34 HEREFORD www.hereford.lug.org.uk/	60 NORTHANTS www.northants.lug.org.uk/
9 BLACKBURN matt@consultmatt.co.uk/	35 HERTFORDSHIRE www.herts.lug.org.uk/	61 NOTTINGHAM www.nottingham.lug.org.uk/
10 BLACKPOOL www.blackpool.lug.org.uk/	36 HERTFORDSHIRE (EAST) www.earlug.org.uk/	62 NUNEATON www.nuneaton.lug.org.uk/
11 BORDERS (ENGLAND/SCOTLAND) welby@mcrobers.com.kg	37 HUDDERSFIELD www.hud.lug.org.uk/	63 ORMSKIRK www.ormskirk.lug.org.uk/
12 BRIGHTON www.brighton.lug.org.uk/	38 HULL www.hull.lug.org.uk/	64 ORPINGTON www.orpington.lug.org.uk/
13 BRISTOL & BATH www.bristol.lug.org.uk/	39 ISLE OF MAN www.iom.lug.org.uk/	65 OXFORD www.oxford.lug.org.uk/
14 BRIXTON blug@communitytechnology.org.uk/	40 ISLE OF WIGHT www.iow.lug.org.uk/	66 OXFORD BROOKES UNIVERSITY www.obaoss.lug.org.uk/
15 CAMBRIDGE mark@taurine.demon.co.uk/	41 KENT www.kent.lug.org.uk/	67 PETERBOROUGH www.peterboro.lug.org.uk/
16 CHESHIRE (SOUTH) www.sc.lug.org.uk/	42 LANCASTER www.lancasterlug.org.uk/	68 PRESTON & LANCASHIRE www.preston.lug.org.uk/
17 CHESTERFIELD & DISTRICT www.spirelug.org.uk/	43 LEAMINGTON SPA ian@siim.info	69 RUGBY www.rugby.lug.org.uk/
18 CHICHESTER HARBOUR – MARITIME GEEKS hugh@mjr.org	44 LEICESTER www.leicester.lug.org.uk/	70 SALFORD (UNIVERSITY) http://linsoc.ussu.salford.ac.uk/
19 CLEVELAND OPEN SOURCE GROUP www.cleveland.lug.org.uk/	45 LINCOLNSHIRE www.lincs.lug.org.uk/	71 SCARBOROUGH www.scarborough.lug.org.uk/
20 COVENTRY www.covlug.org.uk/	46 LONDON (GREATER - GLLUG) www.gllug.org.uk/	72 SCOTTISH www.scottish.lug.org.uk/
21 CUMBRIA www.cumbria.lug.org.uk/	47 LONDON (LONIX) www.lonix.org.uk/	73 SHEFFIELD www.sheflug.org.uk/
22 DERBYSHIRE (SOUTH) www.sderby.lug.org.uk/	48 LONDON (NORTH-WEST) www.northwestlondon.lug.org.uk/	74 SHEFFIELD LINUX & BSD USER GROUP www.slugbug.org.uk/
23 DERRY www.derry.lug.org.uk/	49 LONDON (SOUTH) www.sl.lug.org.uk/	75 SHETLAND www.shetland.lug.org.uk/
24 DEVON & CORNWALL www.dclug.org.uk/	50 LOUGHBOROUGH www.loughborough.lug.org.uk/	76 SHROPSHIRE www.shropshire.lug.org.uk/
25 DONCASTER & SCUNTHORPE www.doncaster.lug.org.uk/	51 LUTON & ST ALBANS www.lust.lug.org.uk/	77 SOUTH WEST www.southwest.lug.org.uk/
26 DORSET www.dorset.lug.org.uk/	52 MALVERN www.malvern.lug.org.uk/	78 SOUTHEND-ON-SEA derek@techsys.co.uk



LINUX USER GROUPS



79 FIFE

www.standrews.lug.org.uk/

80 STAFFORDSHIRE (NORTH)

www.staffslug.org.uk/

81 STAFFORDSHIRE (SOUTH)

ozzy@cowmob.nu

82 SURREY

www.surreylug.org.uk/

83 SUSSEX

www.sussex.lug.org.uk/

84 SWINDON

nick.trueman@ntlworld.com

85 THANET (MARGATE)

patrick@patricknet.net

86 TRURO SCHOOL

tim@rapidgn.com

87 TYNESIDE (NEWCASTLE)

www.tyneside.lug.org.uk/

88 WALES (NORTH)

www.northwales.lug.org.uk/

89 WALES (SOUTH)

www.swlug.org.uk/

90 WALES (WEST)

www.westwales.lug.org.uk/

91 WALTON ON THAMES

rael@freeuk.com

92 WEST LOTHIAN

ben@150bpm.co.uk

93 WEST OF SCOTLAND

woslug@uk2.net

94 WIGAN AND ST. HELENS

www.all-the-johnsons.co.uk/wishlug/

95 WILTSHIRE

www.wiltshire.lug.org.uk/

96 WOLVES

www.wolveslug.org.uk/

97 WORCESTERSHIRE

www.worcs.lug.org.uk/

98 YORK

www.york.lug.org.uk/

99 EAST YORKSHIRE

sharkonline@whatemail.com

100 WEST YORKSHIRE LUG

www.wylug.org.uk/

OTHERS

BENELUX

www.b-lux.org/

BUSINESS LUG

www.businesslug.org/

CHANNEL ISLANDS

myles.hartley@db.com

FUENGIROLA GNU/LUG

www.fuengirola.lug.org.uk/

SCHOOLS

www.schools.lug.org.uk/

STUDENTS

www.students.lug.org.uk/

TR LINUX

www.trlinux.lug.org.uk/

UK UNIX USER GROUP (UKUUG)

www.ukuug.org/

YOUNG LINUX USER GROUP

www.young.lug.org.uk/

Infopoint: NCM fair report



Thanks to **Jono Bacon** and Northern Computer Markets, LUGs can get a free stand at computer fairs to evangelise Linux.

Jono Bacon writes: Over the last few issues, we've been promoting the Infopoint project: the premise is to provide a means for volunteers to go to computer fairs and advocate the use of Linux. So far, there have been fairs in Wolverhampton, Birmingham, Leeds, Hull, Bradford, Manchester and other places. Northern Computer Markets (NCM), one of the UK's most popular computer fair organisers, offered to help kick off the Infopoint project by donating free tables at some of its fairs to demonstrate Linux and give out info and software. Here is an account of one of these Infopoint events.

The place was Dunstall Racecourse in Wolverhampton, and the date was Sunday 13 June 2004. I had put the word out about the Infopoint to the Wolves LUG and we had Ron Wellsted,

Dave and Kat Goodwin, Adam Sweet and Mo Awkati there to help. This was the second Infopoint that I had run; the first involved Matthew Revell and myself. Everyone at Wolves LUG had worked hard to produce posters, LUG flyers, free Knoppix CDs, and lent laptops that we used to demonstrate Linux. From the start, we had a steady stream of interest and ran out of Knoppix CDs and LUG flyers quite early in the day!

When you do an Infopoint, you will naturally need to prepare: making notes of things you need to say and do is a good place to start. It's important to fully explain how Knoppix will *not* touch a user's Windows hard disk: this should be explained very clearly. The availability of a local LUG was of great interest for many of the visitors: many of them found a LUG to be a great support

mechanism if things go wrong or they have questions. It was important to have the LUG website URL and some details printed out on a small flyer to take away.

Possibly the most heartening aspect of running an Infopoint is just how much of a measurable impact you are making in Linux advocacy in your area. Over the course of our five hours at the show, we had over 100 interested visitors, and more than 350 people stopped by the stall with at least a casual curiosity in what we were doing. Even if people don't come and talk to us at the show, the Infopoint concept is great at spreading the 'Linux' word far and wide – this develops brand awareness that can be built upon at a later date.

The Infopoint project has proved to be a fun and satisfying experience so far for the various volunteers involved,

Linux User Group information

Send your LUG-related ideas, event details, criticisms, etc to: **LUGS! LXF**, 30 Monmouth Street, Bath, BA1 2BW or email: lxf.lugs@futurenet.co.uk and spreadingtheword@jonobacon.org. Please note that we ideally need at least six weeks' notice to publicise events in LXF.

but we *still* need your help. We always welcome more volunteers for a huge range of fairs – not just NCM's ones – that we now have access to. Head over to www.jonobacon.org/infopoint/ to join the mailing list, see the dates of Infopoints that need volunteers, and get involved in spreading the word! **LXF**



Freedom

Protect *your rights* to use, copy, modify, and redistribute free software. Become an Associate Member of the *Free Software Foundation*.

Other benefits:

- ✓ E-mail forwarding
- ✓ Bootable membership card
- ✓ 20% discount on merchandise
- ✓ Biannual FSF Bulletin
- ✓ Invitation to annual meeting

Join today at <http://member.fsf.org> or e-mail membership@fsf.org for more information.

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, Mac OS, 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's possible to mount the images and do a network install, or even a local install from a disk partition. The methods often vary between distributions, so check vendors websites for more info. [LXF](http://www.linuxformat.co.uk)

Coverdisc



Neil Bothwick is your guide through the wonders of this month's jam-packed *Linux Format* CD. There's Slackware info here pertinent to DVD readers, too!

On the CD



Wherever you see this logo it means there's related stuff on the CD

ESSENTIAL INFO

On page 111 we have grouped together essential info on the different types of packages on your coverdiscs – along with instructions for installing source packages.

IMPORTANT NOTICE

Before you even put the CD or DVD in your drive, please make sure you read, understand and agree to the following:

The *Linux Format* CD is thoroughly tested for all known viruses, and is independently certified virus-free before duplication. We recommend that you always run a reliable and up-to-date virus-checker on ANY new software. While every care is taken in the selection, testing and installation of CD software, Future Publishing can accept no responsibility for disruption and/or loss to your data or your computer system which may occur while using this disc, the programs or the data on it. You are strongly advised to have up-to-date, verified backups of all important files. Please read individual licences for usage terms.

READ ME FIRST

We have reached the time of the year when all the major distros are releasing new versions. Mandrake released

10.0 official recently, and last month's discs included the new Fedora Core 2. This month, we have two distros on the discs: Slackware has just released version 10.0, and


we have the full version on two CDs and the DVD; and SUSE has made an /ISO image of the Personal edition of 9.1 available under the GPL, which we have on the DVD.

DISTROS SLACKWARE LINUX

Many Linux distributions are derived – at least in part – from other older distributions. Mandrake was originally based on Red Hat, Knoppix has Debian at its heart and so on. If some distributions can be said to have fathered others, then Slackware must be the grandfather of Linux distros. Slackware was the first successful attempt to package up a complete operating system based on the Linux kernel. It was so successful that it is still going strong, more than ten years after its initial release.

Slackware has a reputation for being lean and stable, and sticking more closely to the original kernel and packages it includes. It certainly doesn't include heavily patched kernels, like most of the 'big-name' boxed distros. Instead, it gives you Linux, and the other programs, as their authors intended. It requires a little more effort to install than SUSE or Mandrake, but many people prefer this for the control it gives. This doesn't mean that installation is complex or difficult, just that it requires a little more attention and input from you, the user. Follow the steps outlined below to install and setup Slackware 10.0, from either the CDs or the DVD.

To start the installation, boot from the DVD or first CD. If you have any difficulties with this, see the box titled *How do I boot from this disc?* over the page. There are a number of options



Using *cfdisk* to create the partition ready to receive Slackware.

you can type at the boot prompt to alter the way the installer. Press **F2** to see the choices, but most people will only need to press **Enter** to proceed. When you see the login prompt, press **Enter** again.

Creating space

The first step is to create some partitions to hold Slackware. This is not handled by the installer, you need to use *cfdisk*. See the box that's titled *Partitioning your hard disk* on page 108 for a quick explanation of why you need to create partitions. The Slackware installation disc contains no software for resizing a partition. If you need to shrink an existing partition to make space for Slackware, you should do this *before* booting from the disc. If you already have Linux installed, or have a live CD like Knoppix, you can do this with *QT Parted*. If you want to do it from Windows, use a program

like *Partition Magic*. If you use a Windows program, it is best if you do not try to create the Linux partitions in Windows, this can sometimes cause problems. Shrink your existing partition to make enough room, but leave the space unallocated.

To create your partitions, type **cfdisk** after logging in, move the highlight bar over the empty space on your disk and select 'New' to create a new partition. It is best to create logical partitions when given the choice, or you could run out of available partitions; type in the size and press **Enter**. After creating your swap partition, highlight it, select 'Type' and change the partition type to **82, Linux swap**. When you have finished creating the partitions, select 'Write' followed by 'Quit'. You should reboot after exiting *cfdisk*, by typing **reboot** or pressing **Ctrl-Alt-Del**, to make sure the kernel reads the new partition layout.



COVERDISC CD



You should take a few minutes to read the online help before starting with your installation.

Running the installer

When you log in again, type **setup** to begin installation. If this is your first time installing Slackware, you should read the installation help file, the first option on the setup menu. The next option is probably unnecessary, as you have already been given the chance to set your key map, so select the ADDSWAP option. Each section leads into the next, so you won't return to this menu until the initial installation is done. Setup will detect and configure your swap partition, then ask you to pick the root (/) partition. When formatting the partitions, the default options are fine. By all means, change

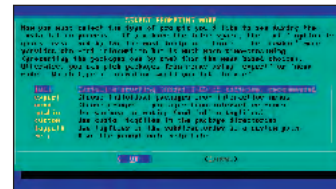
them if you prefer, but if in doubt, take the preselected choices. Once you have formatted the swap and root partitions, you will have the chance to set up any others you may want, such as /home. If you have a Windows partition, you will be given the option of having this added to /etc/fstab. This will not affect your running of Windows, but it will make it possible to read files from this partition in Linux, making it easy to share data between the two operating systems.

When asked for an installation source, pick the CD/DVD option (of course, that is why you are reading this), and let setup detect the disc automatically. Now we come to the real meat of the installation process: choosing the package groups you want installed. The default is to install most packages, which will require around 3GB of disk space. Alternatively, you can prune the list to install only what you need. The next page lets you pick how they are installed. Various options let you choose from menus, be asked during installation or simply to install

the lot automatically. The last option is certainly the easiest. Installing the packages can take a while. DVD owners can put the kettle on now, but CD users will need to be around to change discs when asked.

Final steps

When the packages are installed, you need to choose a kernel. This is usually a case of accepting the default offers, unless you had to make specific choices when booting the installation disc. You should create a boot floppy disk. this will enable you to boot your new Slackware Linux system if anything goes wrong with the bootloader setup. When installing the *LILO* bootloader, a couple of steps later on, you should normally try the 'simple' option first. Despite the warnings it gives later on, most people will want to install *LILO* to the Master Boot Record. If you choose one of the alternatives, make sure you have a boot floppy, or your system may not start. Network configuration is usually a matter of following the prompts. You



Slackware's installer offer more control than most over the installation of packages, although most people will prefer to opt for the fully automatic choice.

shouldn't need to change the startup services, with the possible exception of disabling the PCMCIA service. It was enabled by default when installing on a desktop with no PCMCIA slots.

The remaining questions are fairly trivial except for setting a root password. It is very important that you remember to set a password for the administrative account, even if you are the only person using the computer. You will be able to set up a normal account for daily use later. That completes the initial installation, so you may exit the installer, reboot and remove the CD.

CD CONTENTS AT A GLANCE

Magazine

Beginners
HotPicks
KDevelop
RoundUp
SDLProgramming
TheGimp

Programs covered in *LXF's* Beginners' tutorial
All the programs covered in this month's *HotPicks* section
The various tools used in our *KDevelop* tutorial
The programs covered in this month's *RoundUp*
Source code for *Trout Wars*
The GNU Image Manipulation Program

Desktop

Jreepad
Stellarium
TheGNUPrivacyGuard
Webmin
Xterm

Personal freeform database software
A 3D astronomical sky renderer
A PGP replacement tool
A Web-based interface for Unix system administration
A terminal emulator for the X Window System

Development

DrPython
edtFTP
KDevelop
Steelme
Zinc

A powerful cross-platform IDE for Python
A comprehensive FTP client library written in Java
An integrated development environment for Unix/X11
A drop-in theme manager for Java Swing programs
A graphical toolkit extending the Tk canvas

Distros

Slackware

The latest release of Slackware Linux

Games

EmpireServer
GCompris

A strategy war game
A complete educational suite for children from 2 to 10

Graphics

AlbumCoverArtDownloader
Digikam
GeckoMultimediaSystem

Downloads album cover images semi-automatically
A digital photo management application for KDE
A replacement for your VCR, DVD, MP3, and CD player

Internet

Evolution
MozillaFirefox
MozillaThunderbird
TorrentFlux

GNOME mail client and PIM
A Mozilla-based browser
A total redesign of the Mozilla mail component
A PHP Torrent client and manager

Mobile

AccessPointUtilities
gnuPod
Ndiswrapper

Configure and monitor a Wireless Access Point
A collection of tools for accessing the iPod under UNIX
A kernel module to load windows WLAN drivers in Linux

PBButtonsd
xmitVR3

An Apple PowerBook hardware control daemon
An Electronic Program Guide for the Agenda VR3 PDA

Office

OpenOffice.org
Scribus

An Open Source version of *StarOffice*
A DTP-Program for Linux using the Qt-Library

Sound

Kamix
Mixxx
Rosegarden

A mixer for KDE/ALSA
Digital DJ software
An audio and MIDI sequencer and score editor

System


ClamAntiVirus
ConKolivasKernelPatchset
Deborphan
GRUB
OS-SIM

An anti-virus utility for Unix
A Linux -ck series-based patchset
A Debian orphaned library finder
GRand Unified Bootloader
Security monitor to manage security and network events

Essentials

Allegro
ALSA
Avifile
CheckInstall
CSV
GLib
glibc
GTK
gtkmm
Guile
Jigdo
Kernel
lesstif
libESMTP
libmcrypt
Libsigc
libstdC++3
libXML
Mesa
ncurses
OggVorbis
RAWRITE
SDL
SmartBootManager
SVGAlib

Multi-platform game library
An alternative implementation of Linux sound support
Library to read and write compressed AVI files
A make install installations tracker
Comma separated index files of the coverdiscs
The GLib library of C routines
The C library used in the GNU system
A library for creating graphical user interfaces
A C++ interface for the popular GUI library *GTK+*
An embeddable library implementation of Scheme
Ease the distribution of very large files over the Internet
The latest kernel source and patches
LGPLEd re-implementation of *Motif*
A library for posting Electronic Mail
A library to access various encryption algorithms
A callback framework for C++
The GNU Standard C++ Library
A library for manipulating XML and HTML resources
3-D graphics library which uses the OpenGL API
Text-based interface creation library
Open, professional audio encoding and streaming technology
Write images to floppy disk with Windows
Portable low-level access for multimedia
OS independent and full-featured boot manager
Provides VGA and SVGA modes in a console



Getting Slackware to start with a desktop is a matter of changing one number in one file.

Where's the desktop?

When you reboot, all you will get is a login prompt. don't panic, there is a desktop available, there are just a couple more steps before you can use it. Type **root** at the login prompt, followed by the password you set during installation. Now you need to create a normal user, so type **useradd -m myusername** replacing **myusername** with whatever user name you pick for yourself. Give a password when asked and then type 'logout' (or press **Control-D**) to go back to the login prompt. Log in with your user name and password and type **startx** to load the desktop. Once you see that the desktop loads correctly, you will probably want it to start automatically each time. To do this, you need to edit a file, as the root user. You don't have to log out and in again, you can do it from your normal desktop. This example assumes you

have selected the default KDE desktop, but the principle is similar whichever you use. Open a terminal by clicking the screen icon, fourth from the left on the taskbar. Type **su** to become the root user, it will ask for the password. Then type

```
pico /etc/inittab
```

Find the line that reads

```
id:3:initdefault:
```

and change the **3** to **4**. Press **Ctrl-X** to save and exit and the job is done.

The next time you boot, a graphical prompt will ask for your password and the desktop will load. The Sessions item at the bottom of this login screen allows you to choose from any of the different desktops installed on your computer. These range from the powerful KDE and GNOME, with plenty of options but heavier on system resources, to lightweight window managers that are more suitable for lower-powered computers. Feel free to try them all and use whatever you prefer. Choice is one of the key aspects of Linux.

GRAPHICS DIGIKAM

Digital camera support in Linux is now very good, with the majority of cameras supported. KDE users can access their camera just by typing **camera:/** in the *Konqueror* location bar. While this gives quick access to your images while on the camera, and makes it easy to copy them to your computer, it doesn't help so much


ROOT VS THE NORMAL USER

Root is for administration ONLY

You set a password for the root user during installation, and created another user account later. You may be wondering what the difference is, and why you need two different users, especially if only you use your system.

Linux uses file permissions that set who owns each file or directory and who can read or write to it. Each user has their own directory, the home directory created by the **useradd** command. Within this directory, a user can do as they please, creating and deleting files to their heart's content. Outside of their home directory, each user is normally allowed to read files and run programs, but not change or delete them. This provides safety and security. You cannot accidentally delete files you shouldn't, and even executing a malicious email attachment won't harm the system, because you don't have permission.

When you really do need admin access, say to install a program, you can open a terminal and type **su** to become root temporarily. Some graphical admin tools also need root access. If so, they will ask you for the password, but drop back to normal operation as soon as possible. At no time should you need to log into the complete desktop as root, doing so leaves yourself open to risk.



Setting a root password, and only using it when you need to, is vital for security and safety.

when you have been using the camera for a while and find yourself drowning in a mess of poorly organised pictures. *Digikam* is a KDE-based solution to this, providing all the features most people need for basic photo organisation, and some advanced features through its plugins.

Digikam is supplied as two tarballs, one for the main package and a separate set of plugins. Normally you install both, the main package first, using the standard

./configure && make && make install procedure. *Digikam* manages photos both on the camera and in hard disk, organising the latter into albums. It makes use of the EXIF information embedded in the files produced by most cameras to show extra information.

Albums are organised into collections and photos can be copied or moved between albums with drag-and-drop. Albums can be exported to HTML for web publishing, or archived to CD. The plugins increase the range of possibilities, including creating a DVD slide show. Other possibilities include retouching images in *The GIMP*, batch processing, including renaming and resizing, ideal for web galleries; and a print wizard for hard copies of the best photos.

Digikam is not limited to importing images from a camera. You can import from a disk drive or card reader, which is often faster than downloading from the camera, and certainly easier on the camera's batteries. You can also import images from a scanner, so your old film photos can join the collection. **LXF**

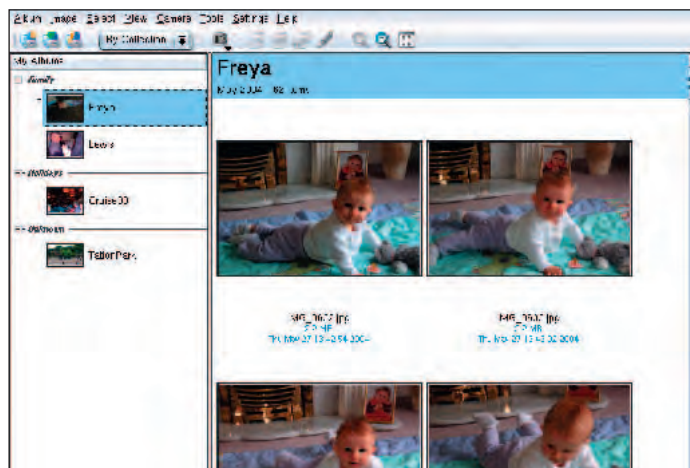
HOW DO I BOOT FROM THIS DISC?

Handling an outright refusal

Although many of *Linux Format's* cover discs – including this month's – are made to be bootable, some people do have problems booting from them. If this affects you, the first thing to check is that your computer is set to boot from CD (this includes DVDs) before the hard disk. Your computer's or motherboard's manual should give details on how to do this, it usually involves pressing a key (often **Del** or **F2**) during startup to access a menu. Watch the messages when you switch on the computer, it will tell you which key to press.

Some BIOSes include a boot menu that allows you to select the boot device each time by holding down a key. This will sometimes work when a straight boot from CD fails, because it gives the system longer to read and recognise the disc.

If your computer is correctly set up but still refuses to boot from these discs, it is most likely that there is an incompatibility between your BIOS and the *isolinux* system used by many CDs. The simple solution to this is to use a boot floppy like *Smart Boot Manager*. This is in the *Essentials* directory of the discs each month. Copy the image from *bootmgr.dsk* to a floppy disk, using **dd** from Linux or **rawrite** from Windows. Booting from this disk gives a menu that should enable you to select and boot from the CD or DVD. In case your computer has no floppy drive, we have made a CD ISO image, called *bootmgr.iso*. Burn this to a CD-R and then boot from that. When you see the boot menu, eject the CD and replace it with the cover disc before selecting the relevant device (usually **CD0**).



Transfer, organise, display and print your photographs with *Digikam*.

Coverdisc



Neil Bothwick is your guide through the wonders of this month's jam-packed DVD. Never has there been a better time to upgrade your old CD-reader/writer!

For the first time in some years, we are able to bring you a current release of SUSE Linux.

This is not a live CD evaluation version, but the latest release of SUSE Linux Personal, 9.1. SUSE is a popular request on the *Linux Format* forums, now we include the new version on the coverdiscs for our DVD readers.

The DVD boots to the Slackware installer, so we have included SUSE as an ISO image, in the Distros/SUSE directory of the DVD. This should be written to a CD-R using any of the popular CD-writing programs for Linux or Windows. See the *Essential Disc Info* on page 111 for more information.

Once you have written the disc, put it in the drive, boot your computer and select 'Installation' from the menu. If




Like SUSE Personal edition? You should consider upgrading to the Pro version for full satisfaction.



The recently open-sourced YaST is the heart of SUSE. It handles the installation and is the program you run whenever you want to configure your system or add programs.

you have problems starting things off, see the box entitled *How Do I Boot From This Disc?* on page 106. If you boot from the disc but your computer stops before the installer starts, there are a couple of things you can do.



The default is to install everything. You can change that here.

Press F2 when it tells you to, this shows the boot messages, which may tell you where the problem lies. Try the options for 'Installation - ACPI Disabled' and 'Installation - Safe Settings' from the initial menu, although most people will have no such problems. You can also change



Wherever you see this logo it means there's related stuff on the DVD

IMPORTANT NOTICE

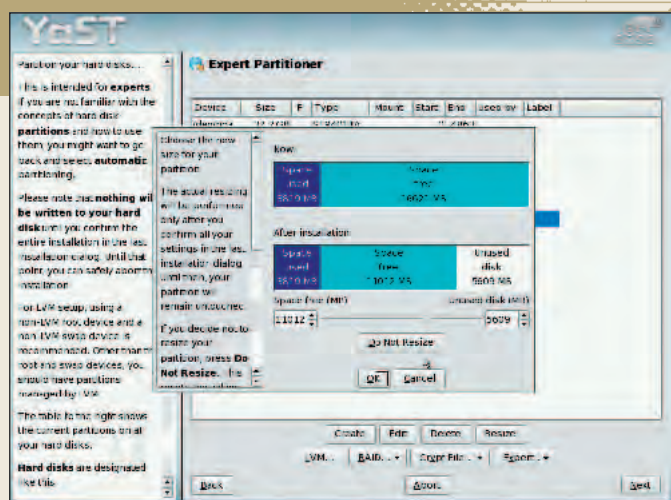
Before you even put the DVD in your drive, please make sure you read, understand and agree to the following: The *Linux Format* DVD is thoroughly tested for all known viruses, and is independently certified virus-free before duplication. We recommend that you always run a reliable and up-to-date virus-checker 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.

PARTITIONING YOUR HARD DISK

A hard disk is divided into partitions, each of which is accessed as a separate data storage area. A standard Windows setup has a single partition occupying the whole drive, so you may not have even noticed that is was partitioned before, but adding a second operating system will require that each has its own portion of the disk. Unlike Windows, Linux installations generally use two or more partitions, the exact number and layout depending on what the machine is being used for.

Ask ten Linux users how you should partition your hard disk, and you will probably get ten different answers! The simplest set-up has a swap partition and a root partition. The swap partition should be roughly twice the size of the computer's RAM, but probably not more than 512MB. It is used for temporary storage of data to free up memory. The root partition (/) contains everything else.


You can also have separate partitions for other sections of the filesystem hierarchy, /usr, /var, /home and /usr/local are the most common



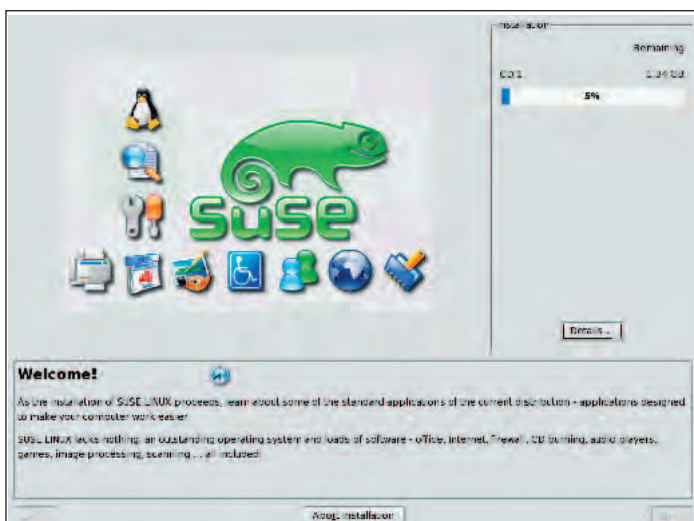
Partitioning your hard disk can be the most nerve-wracking part of any Linux installation, especially if you want to keep an existing OS on the computer, but SUSE's installer takes care of it for you.

candidates for this. However, this normally only complicates the situation. It is generally worth having /home on its own partition. This contains your own data and configuration files, making it a separate partition means you won't lose it if you reinstall. It also means you can share one /home partition between more than one distribution, if you like to experiment

with the different distributions on our cover discs. Some of us also prefer to have /boot on a separate partition, to keep your kernel safe from filesystem corruption. Fedora Core 2's default setup is swap, / and boot, but it is probably best to reduce the size of / and add a /home partition. It makes upgrading or switching distributions much easier.



The main menu, showing a small selection of the software that is now available to you with SUSE.




Installation of the packages can take a while – did you know that the progress bar actually moves more slowly if you watch it?

the display size by pressing **F2** at the initial menu.

The SUSE Personal 9.1 installer is one of the fastest and simplest around. After selecting your installation language, it presents a breakdown of the hardware it has detected and any assumptions it has made. Check through this carefully. You can change

anything by clicking the headline above it or selecting options from the 'Change...' popup gadget at the bottom of the screen. If you already have an operating system installed, check the partitioning section carefully, do not allow it to overwrite your existing partitions, unless that is what you want. See the box titled 'Partitioning your hard disk' for more information on this.

The partitioning section can make reasonable choices for you, or you can use the custom partitioning option where you make all the choices yourself. If you have no unassigned space on your hard disk, this section will let you resize an existing partition to make room. If you want to resize a Windows partition, you should first use *scandisk* and defragment it in Windows to organise the files on there.



YaST summaries what it will do, click on any heading to review or make changes.

CREATING SLACKWARE CDS FROM THE DVD

Subhead xxxxxxxx xxxxxxxxxxxxxx

As is now normal with our bootable DVDs, we have provided a way to create CD ISO images so that you may burn your own CDs for installation on a computer without a DVD drive. Handily, this can be done from Linux or Windows. We have made further improvements to the Linux script, as a result of feedback from users on our web forums.

To build the ISO images in Linux, simply type the following command in a terminal:

```
sh /mnt/cdrom/Distros/Slackware/
mkiso
```

This will create two ISO images in the current directory. If you want to create them somewhere else, give the path as an argument, eg

```
sh /mnt/cdrom/Distros/Slackware/
mkiso /tmp/iso
```

Note that you should not **cd** to the directory on the DVD when running this script. It will not actually fail as it used to with some previous editions of the magazine's coverdiscs, but it will be much slower because the script is not able to open the cache file used to speed up the process.

If you are short of space, it is now possible to create single ISO images with either:

```
sh /mnt/cdrom/Distros/Slackware/
mkiso -d 1
```

or

```
sh /mnt/cdrom/Distros/Slackware/
mkiso -d 2 /tmp/iso
```

If you get an error message along the lines of

```
../Essentials/Jigdo/jigdo-file:
Permission denied
```

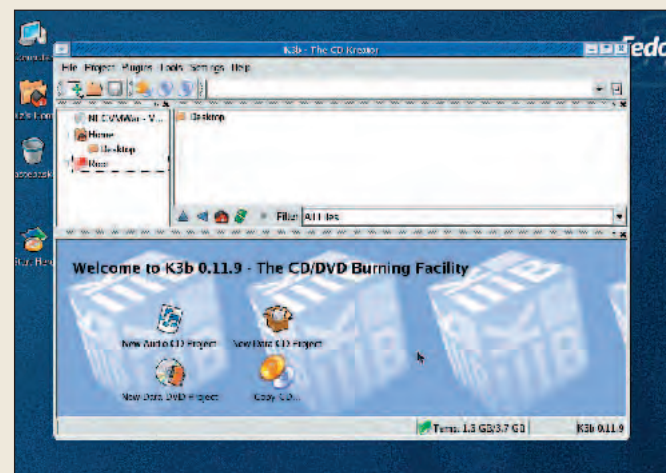
it means your DVD has been mounted with the *noexec* option, which prevents running programs directly from the disc. For security, this is implicit if the DVD is mounted with the user option. In order to fix this, copy the file *Essentials/Jigdo/jigdo-file* from the DVD to somewhere in your path, */usr/local/bin* is a good choice. *mkiso* will now use this in preference to the file on the DVD.

Creating ISO images in Windows

Windows users can create the CDs by double-clicking the *winnkiso* icon or running the script from a MS-DOS prompt. In the latter case, you should change to the Slackware directory of the CD before running the script. The Windows script does not allow you to create single ISOs, but you can specify a destination directory when running it from a DOS prompt. With no argument, or when run from the icon, it puts the ISO images in C:

For example, if your DVD drive is E: and you want to save the ISO images to D:\ISO

```
E:
cd Distro\Slackware
winnkiso D:\ISO
```



After creating ISO images with the *mkiso* script, you can burn them to CD with *K3b*, which is included with Slackware.

The next section on this page shows the software choices. Click on the headline to see details of what is available and what has been pre-selected. As this is the Personal edition, with mainly desktop software included,

everything is selected by default. If space is an issue, you can see how much will be used at the bottom left of the screen, you can deselect anything you feel you will not need. You can always install other packages later by



COVERDISC DVD

◀ running *YaST* from your installed SUSE system. If you remove any packages from the list, press the 'Check Dependencies' button before proceeding, to make sure you have not deselected anything that is needed by another package. Press 'Accept' at the main *YaST* screen to begin the install.

Once the package installation is complete, *YaST* will set up some configuration files, install the bootloader and reboot the computer. This can be a little disconcerting if you

left the computer to get on with package installation and come back to find it rebooted. Don't worry, it didn't crash. After the reboot, there is some more configuration. The first step is to set up the root password. See the box *Root vs The Normal User* on page 106 for more on the differences between the root and normal users. Then you set up your network or modem and have the option to test your Internet connection. There are two good reasons for doing this: not only does it

make sure that your Internet connection works, it also checks SUSE's server for updates. You have the choice of downloading and installing the updates now.

Normal user accounts

Another way: you can run *YaST* again after completing the installation. There were just over 70MB of updates at the time of writing, so it may be better to leave this until you have time to browse them and download the ones

you need. Now it is time to create one or more normal user accounts. You have the opportunity to set up one user account with auto-login. This means that the system will boot directly into that user's desktop, skipping the login screen. This is slightly less secure, but a lot more convenient.

That's it! Exit the installer and your new SUSE desktop should load up. Press the green button at the left end of the taskbar to open the main menu and start exploring. [LXF](#)

DVD CONTENTS AT A GLANCE

Desktop

BlueLava
A Web/WAP CGI-based x10 interface for home automation
CDRDAO
Disk-At-Once Recording of Audio CD-Rs
coLinuxManager
Launches coLinux as a Windows service
DoubleType
Graphical typeface designer
DupeFinder
Locate, move, rename and delete duplicate files
Duplex
Two-sided printing for one-sided printers
Efax-gtk
Provides a GUI frontend for the efax fax program.
ESEKeyDaemon
A multimedia keyboard (userspace) driver for Linux
FreeGuideTVGuide
A TV guide program written in Java
GkrellmSunClock
A *gkrellm* solar plugin
Gnome-pkgtool
GTK+/GNOME front-end to Slackware package management
gTweakUI
A collection of extra preference dialogs for GNOME 2.0+
Jigdo
Ease the distribution of very large files over the Internet
Keychain
A manager for *ssh-agent*
KLinkStatus
A link validity checker
Klogview
A KDE application for viewing log files in real-time
KSmoothDock
A desktop panel for KDE 3.2 with smooth icon zooming
NmapSi
A Qt-based GUI for nmap
Pipemeter
A utility to show the speed of data moving through pipes
Ratpoison
A window manager that lets you say good-bye to the rodent
Remerge
A browser-based interface to Gentoo's *Portage*
Rzip
A compression program for large files
SlackGrade
Slackware update wizard
TightVNC
An enhanced VNC distribution
Usb-mount
USB storage desktop integration for hotplug
Webminstats
A graphing statics module for *Webmin*
Wmconfig
A menu generation tool for various X window managers
Yumtk
A GUI front-end for *yum*
Zisofs-tools
Tools for creating a compressed ISO9660 filesystems

Development

BlitzJavaSpaces
An Open Source JavaSpaces implementation
DoctorJ
A tool for analysing Java code and documentation
PHP
A high-level scripting language
PHPVersion
Check the PHP version and module requirements of a script
PyChecker
A tool for finding common bugs in Python source code
PyQt
Python bindings for the Qt GUI toolkit

Distros

SUSE
SUSE Linux 9.1 Personal

Games

Childsplay
A suite of educational games for young children
NetWormRace
A network snake game for up to 30 players
WorldOfVoom
A turn-based artillery game

Graphics

Divx2dvd.pl
Convert several video files into an ISO DVD file
Fftv
A viewer and recorder for TV and radio
Lxdvdrip
A command line tool to rip and burn a video DVD
mplayerTV
An *MPlayer* frontend for watching TV
NMM
A multimedia middleware package
Varsha
GUI-based DVD authoring software
Vobps2fix
Tweaks .VOB files created by *dvdauthor* for the PS2
Vobstripper
Strip subtitles and extra audio channels from .VOB files
Vobwalker
A tool for splitting a .VOB file into chapters
White_dune
A graphical VRML97 editor and animation tool
XviD
An MPEG-4 video codec

Help

LDP
A complete mirror of the Linux Documentation Project

Internet

BIABAM
Command-line utility to mail attachments
ClarosDownloader
A Web-based download manager
Coccinella
A *Jabber* client with a whiteboard
Dosage
A Web-comic downloading utility
ELinks
An enhanced version of the *Links* text browser
EpiphanyBrowser
A GNOME Web browser
Galeon
A GNOME Web browser
Gtk-gnutella
A *GTK+* *Gnutella* clone
Kasablanca
KDE FTP client with encryption, bookmarks and queues
Laffer
A Web-based ICQ/MSN/AOL/Yahoo messenger
MozillaQuickstarter
An small utility for providing quick launches of *Mozilla*
News-Archive
A Usenet news archiving package
POPsearch
A customisable search engine for the technical community
Qtella
A *Gnutella* client for Linux
WebCleaner
A filtering HTTP proxy
WWOFFLE
Proxy server for use with dial-up Internet links

Office

EnterpriseGroupwareSystem
A Web-based CRM with advanced functionality
Republic
Gets your old legacy report data into *oocalc* spreadsheets

Server

123tkShop
An e-Business solution for small companies
CoherentMailGateway
A greylisting SMTP/POP3 email system for broadband SOHOs
DaDaBIK
Create add/search/edit PHP forms for a MySQL DB
Dnsmasq
A small caching DNS proxy and DHCP server
DSPAM
A server-side anti-spam agent for UNIX email servers
FUDforum
A templatable forum with i18n support
MasswireMail
SQL POP3 Webmail software
Mod_bt
A BitTorrent Tracker implementation for *Apache 2*
MonkeyHTTPDaemon
An small and powerful Web server for Linux
MySQLADMIN
An administration Web tool for MySQL
mysqlCart
An e-commerce solution
PHPLetterIt
A mailing list manager
phpMyAdmin
Handles the basic administration of *MySQL* over the WWW
Postfix
The *Postfix* MTA
Postgrey
A greylist policy server for *Postfix*
PrestoPoll
A poll system with an admin centre

Sound

Agatha
PHP script for streaming music and playlists with *Apache*
Cymbaline
An intelligent learning MP3 player
Enscribe
Convert JPEG and PNG images into audio watermarks
FindDuplicateMusicFiles
Finds music files that are likely to contain the same music
KDERadioStation
An Internet radio streaming directory frontend
Mp3burn
Write MP3s, OGGs and/or FLACS directly to audio CD
Publimark
A public key steganographic tool for audio
QMBTagger
A Qt-based front-end for tagging music with *MusicBrainz*
SmartTagFix
A utility that fills in missing ID3 tags intelligently
Specimen
A MIDI-controlled audio sampler
Wagtail
A client/server audio encoder

System

AegisVirusScanner
A graphical virus scanner for Linux/Unix systems
BackupPC
Enterprise backup system
CDSshell
Bootable CD or DVD console interface and script interpreter
Etherdum
A very small and efficient Ethernet sniffer
Madam
A tool to create, maintain, and monitor Linux Software RAID
RunningUnixMemoryTest
A tool to check the memory without interrupting the service
Slack-get
A tool like *apt-get* for Slackware
Syslinux
Collection of boot loaders for the Linux operating system
UltimateBootCD
Run floppy-based diagnostic tools from a bootable CD
Wasabi
A log monitoring and notification tool

EDITORIAL

Editor Nick Veitch nick.veitch@futurenet.co.uk
Reviews Editor Paul Hudson paul.hudson@futurenet.co.uk
Production Editor Matt Nailon matt.hewnailon@futurenet.co.uk
Art Editor Julian Jefferson julian.jefferson@futurenet.co.uk

Editorial Contributors Jono Bacon, Andy Channelle, David Cartwright, Michael J Hammel, Mike Saunders, Hoyt Duff, Tom Wilkinson, Richard Drummond, Biagio Lucini, Andy Hudson, David Coulson

COVER CD PRODUCTION

CD Editor Neil Bothwick

ART CONTRIBUTORS

Art assistance Colin Nightingale
Photography James Wilson, Powerstock, Photodisc, Corbis
Illustration Paul Bateman, Chris Winn, Shane Collinge

ADVERTISING SALES

Senior Sales Executive
 Andrew Hudson andrew.hudson@futurenet.co.uk
 (01225) 788135

MARKETING AND PROMOTIONS

Marketing Manager Fiona Tully
Subscriptions Executive Susan Stove

PRODUCTION

Production Co-ordinator Diane Ross
Group Production Manager Clare Tovey
Commercial Buyer Mike Thorne
Advertising Production Co-ordinator Jo Crosby

MANAGEMENT

Publisher Stuart Anderson
Publishing Director Dom Beaven
Managing Director Robert Price

DISTRIBUTION AND CIRCULATION

Trade Marketing Manager Jason Comber
jason.comber@futurenet.co.uk

Distributed by Marketforce (UK) Ltd, 5th Floor, Low Rise Building, Kings Reach Tower, Stamford Street, London, SE1 9LS **Tel** 020 7633 3300

Overseas Distribution by Future Publishing Ltd.
 Tel 01225 442244.

Overseas Licences

International Licensing Director Simon Wear
simon.wear@futurenet.co.uk **Tel** +44(0)1225 822798
Fax +44 (0)1225 788105 **Cell phone** +44(0)777 5641493

Contact Details

Linux Format, 30 Monmouth Street, Bath BA1 2BW
Tel +44 (0)1225 442244 **Email** linuxformat@futurenet.co.uk

Subscriptions and Mail Order

Phone +44 (0)1458 271178. See page 100
Email linuxformat.subs@futurenet.co.uk

Copyright No part of this publication may be reproduced without written permission from our publisher. We assume all letters sent – by email, fax or post – are for publication unless otherwise stated, and reserve the right to edit contributions. All contributions to Linux Format are submitted and accepted on the basis of non-exclusive worldwide licence to publish or license others to do so unless otherwise agreed in advance in writing. 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.

All CD-ROM demos and reader submissions are supplied to us on the assumption they can be incorporated into a future covermounted CD-ROM, unless expressly stated to the contrary.

Disclaimer All tips in this magazine are used at your own risk. We accept no liability for any loss of data or damage to your computer, peripherals or software through the use of any tips or advice.

Printed in the UK by Midway Clark (Holt) and Southern Print
 © Future Publishing Ltd 2004

LINUX is a trademark of Linus Torvalds. GNU/Linux is abbreviated to Linux throughout for brevity. All other trademarks are the property of their respective owners

Future Publishing Ltd. is part of The Future Network plc.

The Future Network produces carefully targeted specialist magazines for people who share a passion. We aim to satisfy that passion by creating titles offering value for money, reliable information, smart buying advice and which are a pleasure to read. Today we publish more than 90 magazines in the UK, US, France and Italy. Over 80 international editions of our magazines are also published in 28 other countries across the world.

The Future Network plc is a public company quoted on the London Stock Exchange (symbol: FNET).

Non-executive Chairman Roger Parry

Chief Executive Greg Ingham

Group Finance Director John Bowman

Tel +44 (0)1225 442244

BATH • LONDON • MILAN •
 NEW YORK • PARIS •
 SAN FRANCISCO • WROCLAW
www.thefuturenetwork.plc.uk



Next month

ISSUE 58 ON SALE THURSDAY 2 SEPTEMBER



COOL HACKS

Whether you are a complete beginner or an old hand, everyone enjoys finding out how to do things better, faster, or just with that extra bit of cool. We show you how to get to the core of your system and hack it for extra features, extra performance, and ultimate kudos from your friends!

FTP clients Roundup

Should transferring files be as simple as drag-and-drop, or do more features make for a better product? We pit the best FTP clients against each other to find out.

On test



VIA EPIA MII motherboard, WineX 4, Mono 1.0, C++ Builder X, MandrakeMove 2, SUSE Linux Server 9, and more!

PLUS

DON'T MISS NEXT MONTH'S LINUX PRO!

Dedicated supplement with pages of real-world Linux advice and case studies for IT professionals. Including storage, security and more.

DON'T MISS YOUR COPY OF LINUX FORMAT!
 Get it delivered to your door every month – subscribe on page 100



The exact contents of future issues are subject to change

LINUXPRO

ONLY
AVAILABLE WITH
LINUX FORMAT
MAGAZINE

FROM THE MAKERS OF LINUX FORMAT

SEPTEMBER 2004

FREE EDUCATION!

Open Source software is the natural solution for the education sector, but few UK schools can actually deploy it. Why? Find out inside, and learn from the schools that have made the switch



PLUS

1&1 chooses Linux in-house

Popular hosts 1&1
re-engineer web
controls using Linux

Novell makes plans for Mono

Why Novell and
Ximian are betting
on a C# future

Embedded Linux and Wind River

The big name in little
devices now offers
Linux support

PRACTICAL LINUX SOLUTIONS FOR I.T. PROFESSIONALS

Welcome

TWENTY PAGES OF REAL-WORLD LINUX FOR IT PROFESSIONALS

The generation of people who went to school when 'Computer Studies' started appearing as an option will probably remember well the endless hours of theory, flowcharts and pseudo-code programming that these classes entailed. You were lucky if you actually sat in front of a computer in any given session.

Of course, all this has changed: now computers are ubiquitous, even in primary schools. But the way IT is taught has also changed. Computers are a tool – like a calculator – and students aren't taught how to program them anymore (it isn't even part of the current GCSE), but how to *use* them. Of course, this almost certainly means how to use various applications running on Windows. Take away *MS Word*, and many pupils (and teachers) would be lost. There probably isn't really any need for students not interested in computers to learn any more, but shouldn't IT training include some elements of programming? Understanding how computers work is surely more important than knowing what the shortcut key for inserting a hyperlink is.

This reminds me somewhat of studying Engineering at university. There was a chap who had a terrible memory for the gazillion formulae used in solid mechanics, but a thorough understanding of the theory (and a good grasp of Maths). In exams, he would derive the formula from first principals before using it. Now, he never did terribly well, because he took so long to answer all the questions, but I would rather walk across a bridge that he had designed than an effort from anyone else in my year. Anyway, I digress – find out more about the small successes for Linux in schools this issue. As always, I welcome your comments and ideas at the email address below.

Nick Veitch Editor
nick.veitch@futurenet.co.uk



OF COURSE, IN
REALITY THIS MEANS
LEARNING HOW TO
USE WINDOWS.
TAKE AWAY
MS WORD,
AND MOST PUPILS
(AND TEACHERS)
WOULD BE LOST

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

Stuart Anderton

Publishing Director

Dom Beavan

Website

www.linuxformat.co.uk

Photography

Corbis, Powerstock

ADVERTISING

Senior Sales Executive

Andrew Hudson

andrew.hudson@futurenet.co.uk

(01225) 442244 Ext 5315

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



Published by



30 Monmouth Street, Bath BA1 2BW
Tel +44(0) 1225 442244

CONTENTS

1&1 DEVELOPS

Managing websites the easy way –
1&1's development tale **p04**

FREE EDUCATION

No dark sarcasm in the classroom
– some cheap and educational
Linux boxes would be handy **p06**

WIND RIVER WORKBENCH


Market-leading embedded
developer makes the switch **p16**

NOVELL'S MONO PLANS

Miguel de Icaza on why C# is the
future for Ximian and Novell **p18**



04



06

THERE CAN BE ONLY ONE...

Controlling your websites when they aren't hosted on your own web server is no easy task, which is why 1&1 created its new Control Panel. **PAUL HUDSON** grabs the helm...

Although Linux, Apache, MySQL, and PHP are all available under an Open Source licence, running a website isn't as easy as putting the four together and waiting for the visitors to arrive. Instead, there's a lot of work to be done in choosing the right server, the right networking hardware, configuring and optimising the software, and then making sure it stays up 99.999 per cent of the time.

Sadly, finding a web host with the services you need and one that you can rely on is almost as hard as running a web server yourself; mostly because once your site is served from elsewhere, you lose most of the control you had over its configuration and operation. Although some companies do offer online graphical user interfaces to help people regain a little of the control of over their sites, they're usually limited, prone to doing precisely the opposite of what you want, and rarely let you change the things you want.

1&1 Internet, having more than 3,500,000 hostnames and 1,700,000 active sites, is the largest web host company in the world, and has recently launched version 5.0 of its Administrator Control Panel, which finally offers the kind of usability and control that users have long wanted. What's more, it was developed entirely on and is hosted on Linux. We spoke to Achim Weiss, 1&1's Chief Technology Officer, about the design decisions that went into development of Control Panel 5.0, why Linux was used, and what we can expect in the future...

LINUX PRO: What motivated you to produce the new control panel for sysadmins?

ACHIM WEISS: The ease-of-use aspect was very important to us. As web-hosting solutions become more and more popular – even with less-experienced users – we wanted to create an unbeatable administration tool that is as easy to use as an office solution.

The control panel was an early masterpiece of our in-house R&D department, with some 25 developers having worked on it over the course of its development. We all thought that the previous version was a great piece of software – compared to other solutions it was outstanding, because it allowed users to comfortably configure all major web hosting settings remote on their web browser. But there were some places where the software didn't give users all the help they would expect.

THE FEATURE LIST

From one centralised web interface, Control Panel 5.0 lets you:

- Manage your email accounts, set up mail forwarding, and add virus scanning to mail boxes.
- Administer your domains and subdomains, and buy new ones.
- Use pre-built tools to add more power to your site, such as chat rooms, order forms, and newsletters.
- Set up access to your site through SSH and FTP.
- Track the running costs for your site and manage your invoices.

The new Control Panel 5.0 is very different to the previous release. We wanted to combine the advantages of an in-house Linux-based control panel with an improved set of commands, making sure to design in as much ease of use as possible. We therefore brought in usability specialists who benchmarked the results of the software developers with peer groups of computer users as the software was developed. The feedback given by the usability specialists helped us to improve the overall design and layout of the web interface.

Even technically experienced users now benefit from this new release, simply because more tasks are automated for them: as a result, they can get their work done faster. Consequently, they spend less time thinking about how to operate their server, and spend more time unleashing their creativity. We benefit too – thanks to the software being much easier to use, we get fewer support calls from users.

LXP: What technical reasons were there that motivated the upgrade?

AW: The main goal was to consolidate our different developments from the past so that we could implement new tools faster in the future.

1&1 has a broad range of products which all use a web-based control panel. As all of these products share a lot of functionality at the back-end, it's important for us to share as much code as possible by working from a unified code-base that encompasses the functionality of all control panels. This ensures that we can quickly implement new processes and developments, like new domain names such as .eu or improved virus protection offers.

We achieved this by using our own Java-based framework for web applications called *pustefix*, which we also released as Free Software under the LGPL (hosted at <http://pustefix.sourceforge.net>). This framework is similar in spirit to *Struts* or *Cocoon*. One of the key advantages is that it allows us to completely separate the presentation layer from the back-end business logic. This way it becomes possible to re-use the same code with different user interfaces in different languages and different look and feel. *pustefix* also provides a simple CMS that supports the creation of the presentation layer.

LXP: What system did you have in place previously?

AW: Since we started in 1996, the majority of 1&1's websites have been running on Linux since day one. Even the first web front-end was developed on Linux. Most web applications were written in PHP (starting with PHP/FI in 1997). We gradually shifted from the old system to the new *pustefix* base during the last three years.

LXP: What influenced your software decisions?

AW: Our main software technology strategy for medium- and large-sized systems is based around these tenets:

- Java wherever feasible
- Make it scalable
- Lots of small and inexpensive computers

In our opinion, this leads straight to Linux on AMD/Intel.

Our production systems were originally Debian-based. Today, we have modified them in so many different details and optimised the system kernel for the purpose of fast and reliable web hosting, that internally we call the distribution "Schlunix".

Each of our developer chooses their own development environment based upon what they work with the best. We have Red Hat/Fedora, SUSE, Debian, Gentoo and even some "BSD environments here – it's a real mix. As an IDE, we predominantly use *Eclipse*, but there are also people using *Emacs/JDE* or *Vi*. There are some common standards though, such as *CVS*, *Ant*, *Tomcat* and *Apache*.

As already mentioned, we use a Java framework, so most of the web apps and business logic is written in Java, but we also do a lot of system programming in C++ and rapid application development in PHP.

Our requirements for software systems are quite special. Being the world's biggest web hoster (according to Netcraft) means that there are no standard software solutions. In our experience, it proved to be quite successful starting the development with existing Open Source software instead of starting from scratch. We try to feed back our changes to the projects to avoid maintaining special versions. Furthermore, Open Source software gives us the possibility to fix problems ourselves instead of relying on possibly unreliable vendor support.

LXP: How long did it take to produce, from initial concepts to final implementation?

AW: The framework itself has been constantly improved since mid-2000, and comprehends several man years of development. The current web hosting 5.0 Control Panel builds up on code we already had developed for earlier products, so it's impossible to give a meaningful number for this product alone. Just the final finishing for 5.0 kept four developers busy full-time for six weeks.

LXP: What were the key problems you had to solve?

AW: Separating content and business logic is always the main problem. You need to find a good balance between generality and complexity, so the whole chain of developers – from framework to application developers right up to the HTML authors and designers – can understand and work with the system without the need to understand the whole framework in every little detail.



Once we had a system in place that forced us to think of the business layer as an entity independent from the user interface, the problems that remained were just the usual day-to-day problems of application development.

LXP: If you had to do it again, what would you change?

AW: If we had to start all over again, we would research the possibility to run the framework inside of an application server like *JBoss*. Currently, we can use business logic residing inside our application servers, but it would maybe be worthwhile to see what we would gain from a stronger coupling between the front-end's controller logic and the back-end business logic.

LXP: What kind of feedback have you received from your customers about the new system?

AW: Our customers have reacted very positively. They like the new look and feel, which feels friendly and very tidy. We don't have exact figures just yet, but all our support agents confirm that questions about how to use the web interface have been reduced significantly.

LXP: What can we expect to see in future releases?


AW: There is nothing that's good now that can't be done even better in future, and we're still learning from the feedback of our customers. However, I can say for sure that the Control Panel of the future will look and feel even more like an application rather than a simple website. ■■■

Achim Weiss, Chief Technology Officer, 1&1.



"THE MAJORITY OF 1&1'S WEBSITES HAVE BEEN RUNNING ON LINUX SINCE DAY ONE. EVEN THE FIRST WEB FRONT-END WAS DEVELOPED ON LINUX."

ALL PHOTOS: 1&1

COVER FEATURE **EDUCATION**

EDUCATION EDUCATION EDUCATION

FREE EDUCATION!

Concerned parent **RICHARD SMEDLEY** learns what schools are doing to free pupils from inflexible and expensive proprietary software, and outlines what we can all do to get education and learning back into the current regime of office software training.

People that left school twenty years ago with a computer studies 'O' Level may not have spent much time using a real computer, but certainly would have had a firm understanding of them. You knew the principles of the hardware and – more importantly – how to program them. Best of all, computers weren't a magic mystery box, they were a machine – something to be understood. Today's pupils spend a lot of time using computers – they are part of the curriculum for every subject – but those who take the ICT (Information & Communication Technology) GCSE may spend a great deal of time using spreadsheets from a well-known US monopolist, but they gain no real understanding of what is going on inside the magic box. Far worse, put them on a different machine with a different version of their OS and office software (or even another OS and office package), and they'll find they have only been trained – not educated enough to quickly make themselves at home in any ICT environment.

Mark Chapman of Trinity High School in Manchester sums up the current course: *"We don't do computer studies: we teach Information and Communications Technology – teaching children how to integrate that technology into their lives. Programming is no longer part of the curriculum at GCSE level."* Well, our offspring may not all want to be computer programmers, but education for education's sake demands that children have a fair chance of learning logic, problem-solving and analytical skills. Such skills as may be imparted by tasks like programming a computer, say.

A programming GCSE is available, but the problem is that hardly any schools offer it, lacking trained staff. In fact, while

the government throws expensive hardware and cut-price (but ultimately very expensive) Microsoft School Licensing Agreements around, schools struggle to find competent IT technicians who'll work for half the salary of a teacher. But what schools teach is only part of the problem – the closed nature of the software also sits at odds with pedagogical ideas of knowledge-sharing and community. (The MS Schools Licensing Agreement is currently the subject of a UK Office of Fair Trading investigation – see pages 13-14 for more on this.)

In this article, we will examine some of the problems that schools face in getting Free, Libre and Open Source Software (FLOSS) onto their desktops and servers, and look at some of the schools that have overcome the problem. We will also see that many other countries are way ahead of the UK in this area for various reasons. One of our problems lies with centralised control of schools and their curriculum, and restrictions on staff-training. Sadly, the best hope for change lies with a government which still seems hopelessly wedded to expensive proprietary software and unrealistic techno-fixes

TODAY'S PUPILS SPEND A LOT OF TIME USING COMPUTERS AS PART OF THE CURRICULUM FOR EVERY SUBJECT, BUT THEY GAIN NO REAL UNDERSTANDING OF WHAT GOES ON INSIDE THE MAGIC BOX...

(such as 'a laptop for every pupil' – hardly practical when many pupils seem unable to remember to take a pen to class (www.private-eye.co.uk/content/showitem.cfm?issue.1108/section.teacher). This is something that readers can change: when MPs find that enough of their constituents care about an issue, they take notice and start to care too. This article gives details of some of the campaigns that readers may care to support in order that today's pupils can benefit from the freedoms conferred by our favourite OS. >>>

POWYS LEADS THE WAY

Top marks for forward-thinking in rural areas

IF YOU HAD BEEN LOOKING AT THE SITUATION of FLOSS in UK schools three or four years ago, you would have found several individual schools doing their own thing and communicating via the suse-linux-uk-schools mailing list. Yet, you would only have found one LEA embracing FLOSS. Fast-forward to the present and unfortunately the situation is still much the same, with the lead in desktop roll-out now also being taken by Powys.

<http://itteam.powys.gov.uk/home/>

Powys is a large county with a rural population attending many small schools – particularly at the primary level. "All schools (or at least primary schools) within Powys have a Linux box that acts as a dialup server, email gateway and file server", says Sion Pennant, Web Development Team Leader at Powys County Council. "They have an in-house developed administration facility based on PHP and MySQL. These were developed using the National Grid for Learning scheme."

Nick Talbott, ICT Policy & Planning Manager at Powys LEA, and his team rolled out this Slackware Linux-based system seven years ago, and its continued use is a hearty tribute to its robustness. Schools administrate their own boxes using the browser-based front end, with (rarely needed) support available from the LEA. An in-house developed mini-CMS has also been taken up by a number of Powys schools. Martin Williams, head of the LEA IT Support service, has found most of the problems down to trivial failure of cheap hardware, such as CPU fans, while the Free Software has proved most resilient.

Cost-per-school for software licences increases in rural areas, which explains why Powys first

looked at FLOSS. Their continued success can be explained by the unusually high (for an LEA) availability of skilled 'nix technicians – again, a possible by-product of the isolated, rural geography. Talbott also cites "a good relationship between the Council's IT Centre and LEA school's IT team. Many other sources lead us to believe this is not all that common! Our IT Centre has traditionally had 'nix skills." Powys is also involved in the current e-Gov FLOSS trials, and now Powys is investigating a desktop rollout. Maldwyn Pryse of the LEA is in particular interested in Welsh language support in software for schools.

LANGUAGE SUPPORT

OpenOffice.org 1.1.2 has Welsh language support way ahead of what the Welsh Assembly is paying Microsoft to put into the next edition of MS Word; and KDE and GNOME localisation is usable, with completion on the horizon. Add in Cymru (the bootable Welsh distro), and it's increasingly hard to see why every LEA in Wales hasn't chosen to adopt FLOSS already.

So, can Powys' success be replicated in areas without its special circumstances? Wales-based consultancy Ateb has rewritten Powys' solution and rolled it out in neighbouring Ceredigion, and in Worcester. Simon Annetts of Ateb notes that Worcester LEA "is very committed to using Free Software as much as possible, and all but a few of its 200-300 schools have a server appliance installed with our software on it, providing authentication, webmail, file and web proxy/filtering services. The entire central infrastructure is also built on Free Software, mail

and DNS servers, web and caching servers as well as central LDAP authentication to their Content Stream service. Pupils get a single login to both their school server services and also to all externally provided online learning content using just the one account. Account replication between the two systems is handled automatically."

Oldham also has FLOSS solutions for schools' web services. Unix and Linux System Administrator Andrew McCall notes that the LEA is not actively pursuing FLOSS but has no problem implementing it where appropriate: "It's all about what services they need and how they want to implement it, but they are not going to pursue open source software unless they know it exists, that's where admins like me and magazines like Linux Format come in handy." Some cause for optimism then, with successful examples to point to – remember this when you are talking to your children's schools.

« EDUCATION EDUCATION EDUCATION

GETTING THE GNU INTO SCHOOLS

FLOSS adoption depends on the expertise of the implementer and enthusiasm of staff and pupils.

Lets take a look at a few of the schools that are running GNU/Linux at some level. There are probably several hundred schools in the UK running at least one server with Free Software, usually web-caching with *Squid*, or perhaps providing email or file/print sharing. However, a lack of support from most LEAs or Becta, combined with the all-pervasive Microsoft School Agreement (www.microsoft.com/Education/SchoolAgreement.aspx), means that the few establishments attempting to do more than this are of particular interest. The case studies here represent the range of installations, and the problems encountered in getting FLOSS into UK schools.

EAST HULL CITY LEARNING CENTRE (CLC)

East Hull CLC rolled out a thin-client Linux solution on 120 desktops some two years ago, based on the UK-developed answer to LTSP: *Nymph*. Like any large software rollout, it has needed a few tweaks, but generally things have gone well, and staff and management are pleased with the results so far. ICT Director Dougal Gill has *"found the Linux operating environment to be robust and easy-to-use for the limited scope we have so far undertaken. (Principally Star Office, Mozilla and Konqueror web browsers) There was initial problem with running Flash animation, but this is now fixed. RealPlayer files continue to be a problem, though I am led to believe there is a solution in the pipeline."*

More RAM has been added to terminals to make a 'medium client' solution to help with bandwidth congestion problems at peak log-on times, though there are some alternative workarounds currently available from the LTSP community. The CLC is also experimenting with a wireless network solution, for reasons of both safety and cost.

Despite some inevitable complaints about differences between *Star Office* and *MS Office*, especially when it comes to databases, *"teachers took to Linux pretty well"*. The big gain as far as Gill is concerned is in the cost-savings, some £70,000; *"also, students at the centre can purchase fully-licensed copies of Star Office for home use at £2, through Pebble Rock"*. Use of *StarOffice*, or the entirely free *OpenOffice.org*, closes the digital divide, and allows schools to recommend software for homework completion without tacitly encouraging the unauthorised copying of proprietary software beyond the financial means of many homes.

Gill not only intends to take The CLC's FLOSS use further, but is also becoming an increasingly important advocate of FLOSS in education: *"We intend to look at other educational software with a view to 'badging' it as 'Linux-friendly', working with suppliers to increase the range that will run on the Linux platform, but this is still in its early stages."*

The 120 thin-client computers are housed in specially designed desks, and the six Linux servers are rack-mounted


in a large server cabinet, as shown in the picture below. Each server has 1.5GB of memory, an Athlon 1800+ and a 60GB hard disc drive – a relatively inexpensive setup. Each client has a network card with specially programmed Boot ROM – no need for a hard disk drive, CD-ROM or floppy. This left enough extra money to buy 15inch TFT screens. The servers are remotely managed by The Learning Machine.

East Hull CLC does have one Windows 2000 server, and one room with some Windows 2000 workstations: they are there in order for the children to use specific software and legacy applications that have been written for sole use with a Windows operating system.

IF FLOSS IS CHAMPIONED BY AN INDIVIDUAL, WHEN THEY MOVE ON, IT'S NOT JUST TECH SUPPORT THAT GOES, IT'S ALSO THE POLITICAL WILL



EAST HULL CLC: Yes, they're Linux terminals – the screens present the same angle as a book open on the desk, and the kids can't get up to mischief hiding behind monitors!



The simple Linux power for 120 simultaneous desktop users at East Hull CLC.

ICT manager Dave Langcaster speaks of the difficulty in finding adequate training for tech staff that is specific to the types of tasks carried out in schools. Similarly *"NOF training for teachers is targeted at Microsoft software, and teachers generally have experience of Microsoft. Also, teaching materials are all Microsoft and so have to be rewritten."* The CLC has also faced a low-level *"reluctance by teachers and members of the public to embrace Open Source."*

Nevertheless, despite these hurdles, East Hull CLC remains forward-looking and optimistic about the future of FLOSS in education. In addition to looking for funding to develop some of the 'missing' software – *"Internet control, classroom monitoring (such as NetOP), printer management, adequate database with Access-like interface"* – East Hull CLC is *"also looking to allow schools to access thin-client sessions over the Internet."*

MANCHESTER

England's third city is an interesting microcosm of the problems that face education, and the solutions that have been tried – and the same can be said about attempts to bring FLOSS into schools there. 127 schools in the Manchester Education Wide Area Network (MEWAN) have at least one Linux box as proxy, traffic filter or web cache.

Directory services and content management have just been rolled out. As access to these is browser-based, they are available to users on any platform. Schools with CLC status – such as St Paul's in Wythenshawe, Parrs Wood, or Trinity High School – have been able to take things further. Mark Chapman, who heads ICT at Trinity, has been particularly impressed with *Moodle* (see page 13). Browser-based services make management easier, and anything

PHP-based is easy to roll up your sleeves and get working on. Ben Higginbottom, a key Linux advocate among the technical staff, also sings the praises of the network model. Higginbottom has several SUSE Linux servers to admin at Trinity, running DNS, LDAP, DHCP, etc.

Chapman is keen to see diversity in ICT education: *"I don't want children to be calling a word processor Word. I don't want them to be calling a spreadsheet Excel; and I couldn't show them alternatives anymore. At the very least I wanted to show them that it doesn't all have to be produced the same way."* Next year's Year Seven (11-year-olds) will have dual-boot machines – and will use OOo on GNU/Linux. *"If there was a database that I could do mail merge and things like that; and had validation rules, fields, queries in Access-like ways, I could start Year Ten on it,"* continues Chapman, who is unable to introduce FLOSS where it is not ready. Pupils arrive having used MS Access at

primary school, and neither *Total Recall*, nor the integration of *MySQL* and *OpenOffice.org* is ready to replace *Microsoft Access* in schools – yet.

St Paul's uses the versatile thin(ish) client system, *Nymph*, running on SUSE 8.2 – the school has not taken up paid support, so the system has not been updated. The situation is further complicated by a new head that makes no secret of their antipathy to FLOSS. This illustrates the problem of schools with FLOSS championed by an individual, and what happens when that individual leaves – it is not just the understanding of the FLOSS sphere and tech support that goes, but also the political will. The Manchester schools are also plagued by low skills and motivation in a number of the technicians, and staff hostile to changing the software that they know how to use. These are issues that nearly every school will struggle with unless the political will is found to provide training in *nix, and support for FLOSS rollouts. >>>

FOLLOW THE LEY LINES

Glastonbury is much more than a festival...

STEVE LEONARD CLARKE, HEAD OF ICT AT St Dunstan's school, Glastonbury, gave *Linux Pro* the low-down on the successes and failures of the school's comprehensive Free Software roll-out.

STEVE LEONARD CLARKE: In March 2002, we installed one Mandrake 9.1 Enterprise server with 15 thin-clients (Pentium 90s donated by pcsforschools.org). The Learning Machine and Fensystems manage them remotely via *SSH* over the Internet. In January 2003, we installed two more Mandrake servers (load) as well as 15 brand-new (dual-boot) mid-client machines with Windows 98SE on them along with another 15 (dual boot) thin-client recycled Pentium PCs.

LXP: And your successes and failures...

SLC: Our more notable successes include:

- We can run generic office applications and have Internet access on old machines that would otherwise be useless (a budget deficit meant no expansion for us, so we had to come up with other ways to expand at little or no cost).
- Students learn to work on different platforms and have learnt about FLOSS philosophy.
- The PC/pupil ratio has dropped from 1:9 to 1:5, so more students can use technology. We can teach a class of 30 and no student needs to share a PC.
- Staff have seen the importance of GNU/Linux.
- I don't like to sound like I'm blowing my own trumpet, but implementation went relatively smoothly, as I knew what I was doing!
- Winning the support of my senior managers and governors was essential. (It took a while to start with, but was a success in the end).
- The Learning Machine, on the web at www.irlcomputers.com offers excellent support.

Here is what we consider to be the failures:

- The LEA held us up with the installation by six months. At LEA level, there was a defined lack of awareness of the issues involved, and it was difficult to get its support – this issue is still ongoing.
- Because of our restricted budget, we have had to use second-hand machines – some students thought they were getting a 'cheap' deal. Again, implementation had to take this into account and each group was carefully involved with the philosophy behind the strategy at St Dunstan's.
- We are still having problems with the mid-client machines. I want to run *Rosegarden* as well as other GNU/Linux audio/multimedia applications on the new machines, but as yet we have not had time to set this up properly.
- Essentially, money is the main limitation as to what we would like to set up at St Dunstan's – the old network is NT and is still in use by most departments. We could not afford to pay the licence to connect the two networks together, so pupils have two home directories at present. The plan for next year is to remove the NT server altogether and make the backend completely GNU/Linux, and having a mixture of Windows and Linux desktop machines on the network.

LXP: What were the pupil and staff attitudes to these changes?

SLC: Both pupil and staff attitudes have been surprisingly positive, but again I think this common when 'something new' is introduced. I teach most of the ICT in the school, so pupils were able to get their information 'from the horse's mouth'. Staff were able to have proper INSET from me via a presentation of the strategy, as well as getting first-hand experience through training sessions.

LXP: There must have been a fair amount of money saved?

SLC: For the governors and the local press, I put together some figures based upon setting up a new Windows XP network with *MS Office XP* and *Adobe Photoshop 7* (full version), in comparison with setting up a GNU/Linux network with *OpenOffice.org* and *The GIMP*. The figures were quite astounding: on the first Linux network, we saved £9,000 approx; the second expansion saved £25,000 approx. This is on software alone!

LXP: That's great – what about the educational improvements that a FLOSS policy has brought?

SLC: Some are quite difficult to quantify!

- Pupil/PC ratio is better, therefore pupil experience of ICT in action is improved.
- Pupils are directly involved with the notion of Open Source and Free Software – an educationally sound philosophy.
- It has allowed us to begin implementation of the new Key Stage 3 Strategy, while not affecting the use of ICT in other Departments.
- I set up the Linux User Group of Glastonbury (www.lugog.org) in March 2002, which has brought in local interest as well as local GNU/Linux and Unix expertise, which is of mutual benefit to the school.
- Next year, I will be introducing the international grades in office technology certification scheme (www.theingots.org) using *OpenOffice.org*. This will mean that students can get certification at bronze and silver level by the time they leave St Dunstan's (it is a bit like the ECDL).
- By using *the opencd* (www.theopencd.org), pupils can use the same software at home as in school freely, without complicated licensing issues.

« EDUCATION EDUCATION EDUCATION

THE GOVERNMENT AND THE LOBBYISTS

FLOSS adoption quoted by government agency as a success, despite receiving no official support.

Friendships between Prime Ministers and software billionaires aside, the UK government as a whole lags behind its European neighbours in expressing any coherent IT policy for either education or healthcare, let alone use of Free Software. All decisions in UK education are centralised, and Linux is only found where intrepid individuals have stood against the system. So, where do these decisions come from, and what can be done to change them?

Political will at the top is as lacking as understanding among the decision makers of both the technical issues, and of what Free Software is all about. While industry is moving wholesale to GNU/Linux, education decision-makers are throwing every scrap of their expanding budget into proprietary software. The Liberal Democrats have made all the right noises on FLOSS, but the LibDems' voting on software patents in the European parliament undermines their stated position. Backbenchers in the two main parties have asked pertinent questions in the House on FLOSS, but the Department for Education and Skills seems even more strongly wedded to Microsoft and other proprietary solutions than many other departments. If Becta doesn't back FLOSS, how will it reach the majority of schools? And what can we do to encourage change?



BECTA

Becta (the British Educational Communications and Technology Agency) is one of those government agencies that are not directly a part of the government, but have all the powers and more. Becta's purpose, according to Chief Executive Owen Lynch, is to "support the transformation of education through the exploitation and embedding of technology in learning and teaching, in educational organisations, and in developing wider education networks and systems". Computers are to be used in every area of teaching, just as they have found their way into most areas of post-school life. A laudable aim, if carried out with due sensitivity to putting learning first.

Becta is, in its own words, "the Government's key partner in the development and delivery of its Information and Communications Technology (ICT) and e-learning strategy for schools and the learning and skills sector". Which loosely translated, means that if the schools want information on anything computer-related – including what software they should be using – it is to Becta that they have to turn first. And here they find that proprietary software is the way to go; just as several years ago, Becta told them to get rid of their

SCHOOLFORGE UK

www.schoolforge.org.uk/

FOR YEARS, THE LOBBYING EFFORTS OF concerned individuals in the FLOSS scene have languished in dead-end campaigning projects. Now, at last, those who wish to see long overdue change have the twin efforts of AFFS and Schoolforge-UK to rally round. The Association For Free Software (AFFS) sees education as something of a priority and is recruiting active members to its campaign. It has organised the Free Libre and Open Source Software In Education Conference (FLOSSIE) in London, which will become a regular event (moving around the country). Schoolforge-UK is a natural partner for AFFS in this, having a wider remit than purely software, but objectives that nevertheless lie in the same direction.

Richard Rothwell took time to explain Schoolforge-UK's mission to *Linux Pro*:

LINUX PRO: What is the main aim of Schoolforge-UK?

RICHARD ROTHWELL: Schoolforge-UK aims to encourage and support the use of FLOSS in

schools in the UK. We are a young and developing grouping aiming to full together all the interested parties and to lobby for wider use and awareness of the potential of FLOSS. Further, we have a long-term view to do with trying to produce resources that can be build upon and re-used rather than 'supplanted'.

LXP: What has it achieved since its founding?

RR: In the year or so since our inception, we have opened conversations with Becta (www.becta.org.uk), having its Chairman talking at our conference this year. We are liaising with other organisations to expand knowledge and understanding of FLOSS. We are involved in a project to understand the TCO implications of adopting Open Source solutions in schools.

LXP: What does Schoolforge-UK aim to achieve this year?

RR: We have a number of projects in hand at the moment key in these are:

■ Expanding awareness – in particular I will be

addressing the ACITT conference (<http://acitt.digitalbrain.com/acitt/web/events/menu/?verb=view>) in July, demonstrating LTSP and showing the current state of solutions such as *OpenOffice.org*. We hope, in association with OpenAdvantage, to at least have a strong showing at BETT (www.bettshow.co.uk/) this year, and possibly a seminar stream.

■ The development of a FLOSS database reporting and form design tool – at the moment this is the 'gap' that we perceive in the complete FLOSS solution in schools (and SMEs for that matter). We are working with a group of Birmingham universities, organisations and SMEs to gain funding for the writing of such a tool. This would be expected to be a superset of those offered in OOo 2.0, though it would be compatible with and capable of integration with that solution.

■ Further collation and development – as Prof Hargreaves discusses in his Demos publication (www.demos.co.uk/workinglaterally_pdf_me_dia_public.aspx) the Open Source development model is particularly appropriate for education.

British Acorns and move to the Wintel platform. Schools don't have to worry about the expense of software, with £110million in e-learning credits floating around; nor do they have to overly worry in many cases about paying for hardware: there is often money available for this as part of increased government spending. When looking for funds to train staff on *nix though, schools everywhere find that the cupboard is bare.

Four years ago, Becta first began to take an interest in FLOSS (which it refers to purely as 'open source'). Things looked promising as a two-day Open Source conference was planned in the Midlands, and grants were offered for FLOSS projects in schools. Suddenly, and without warning, the plug was pulled on both conference and grants, leaving one large project to trial GNU/Linux on the school desktop without funding. A Becta insider resigned amid rumours of Microsoft disapproval for the initiative, and UK schools continued to be directed down the proprietary route.

So where is Becta now? Well, according to a prepared statement from Dr Stephen Lucey, Executive Director of Educational technology at Becta: *"Becta takes open source software seriously, and understands its potential benefits and potential limitations. Part of our role is to select suppliers and provide procurement advice to learning institutions and Local Education Authorities: minimum specifications for tender and accreditation projects are deliberately made sufficiently open to tender responses from open source providers."*

"We also provide a range of advice on licensing, use of and benefits of using open source, such as 'What is open source software?' and 'Using open source software.' This can be found on our ICT Advice site: www.ictadvice.org.uk" Unlike the main Becta site, putting Linux into the search engine here at least brings up some results. Eight results in fact,



including the above-mentioned 'Using open source software', which give a brief point-by-point summary of some of the advantages of using Free Software (*still* always referred to by Becta as 'open source'). It also mentions thin-client solutions and the Linux Terminal Server Project. Although the links are a little out-of-date, even odd, it's not a bad document – perhaps Becta could revamp it and then send it to schools, instead of hiding it away in its knowledge-base.

Lucey also pointed out that *"several open source initiatives are underway in UK education: for example, Powys LEA has deployed Linux based servers to all schools; Manchester LEA has implemented a central Webmail service based on Linux and gmail; CardiffSchools.net has implemented personal websites and email for schools in the Cardiff management."* Yet somehow he forgot to mention that **Becta played no part in any of these initiatives**. Lucey also claims that Becta recognises FLOSS's role *"in creating more home/school links: finding ways to be more inclusive is one of Becta's priorities, so any potential cost reductions are carefully considered."*

However, he goes on to raise the TCO issue so beloved of Microsoft and its sponsored reports. Becta has started TCO trials in 22 schools using an online TCO tool (results will be available in 2005). Though Becta has not been forthcoming about the details, it also intends this, it says, as a chance to see if FLOSS can meet the specific needs of education. *"Schools have limited technical support staff, which could, potentially, change some of the regular financial arguments for open source"*, continues Lucey. Well Becta, how about getting some of the £100million+ per annum of e-learning credits redirected to staff training?

Feeling generous, one might award Becta a C+, but the comment on its report card would be "must try harder".

LXP: Your objectives speak extensively of FLOSS, yet some observers might be critical of the fact that your aims seem to deprecate software in favour of a broader aim of Open content. What is this all about?

RR: We aim to strike a balance. As I have suggested earlier, central office tools are crucial to some of the delivery of ICT in schools, but the content is in many ways more important. The openness and permanence of HTML is a clear example – whereas I cannot easily use programs written for the BBC micro, even though some of these are excellent but lost resources. A clear example is the splendid BBC Domesday project (www.iconbar.com/news/features/camileon.html) which is all but lost because of its closed nature. We believe that anything produced today should be available for re-use and further development to meet future needs.

In terms of software, we are exploring ways of encouraging and supporting the production of Linux-friendly software – and are considering launching a 'Runs on Linux' branding – possibly in

the same way that the 'Curriculum On-line' mark works. Over the next term, we will be looking at the needs of primary schools and identifying the key software that needs to be ported or produced for their specific needs.

LXP: How does Schoolforge-UK actually work? There is rather a quiet forum at the site, and seems to be no mailing list. So, how do members network to achieve the Schoolforge-UK objectives?

RR: The website and forum are relatively quiet, as we have no intention of attempting to replace the already efficient channels – such as the general suse-schools mailing list (www.suse.co.uk/uk/company/schools/), and specific resources such as the LTSP list (www.ltsp.org/maillinglists.php). There is a mailing list and the Wiki is used in bursts for the production of collaborative documents. We tend to meet at events of shared interest – as we are at the moment completely unfunded! So the AFFS conference (which is being hosted by the

school that I teach at) will be such an event. (www.affs.org.uk/).

LXF: What's the main thing that you're working on at the moment?

RR: Our main aim at present is increasing awareness. We still have to explain what FLOSS is to most people we talk to – and they are unaware of the current viability of an Open desktop solution: I guess we are not alone in this problem! This involves increasing awareness inside government of the issues – running the Government's own e-Government Interoperability Framework (e-gif) website through the W3C HTML validator produces remarkable results (<http://validator.w3.org/check?uri=http%3A%2F%2Fwww.egif.org%2F>). Similarly, the recent eLearning consultation document – to which we responded – was an extremely unpleasant (tied-down forms, etc) MS Word document. In terms of practical application, our lead in the FLOSS database report and form writer should produce tangible gains for education and the wider commercial world.

« EDUCATION EDUCATION EDUCATION

EDUCATIONAL SOFTWARE

Work by pupil-led projects plugs gaps left by UK government's £100million e-learning initiative.

The Free Software community has responded to many schools' needs for all sorts of specialised software. However, in the UK there is little incentive to write software for schools who will not use it. Following the advice of Becta and the LEAs, schools mostly buy proprietary software. In fact, the government's Department for Education and Skills (DfES) spends a staggering £100 million *annually* on proprietary licensing through the e-learning credits initiative. It is hard to see that paying teams of British developers individually to produce all the software needed under FLOSS licences could cost more than a fraction of that sum – yet would bring continual improvements from the community.

Schools and LEAs need a staggeringly diverse amount of niche software: from car-leasing schemes to child protection, through general admin apps to software that actually relates to the curriculum. Most of the time, however, it boils down to just a few core apps – with an office suite and a web browser being the main ones at secondary level.

Primary school applications are a little more specialised. A good start for the needs of younger pupils is *GCompris* (from the French for "I've understood"). This GPLed app, which runs on MS Windows as well as 'nix, is a 'board' (which uses the *GTK Canvas* widget) onto which educational games for 2 to 10-year-olds can be added. Included in the main package are apps from algebra and learning counting with money, through memory games and a simplified Tower of Hanoi, to logic-training and music activities. Part of *GCompris'* popularity is down to the ease with which new boards can be created – you can even involve older pupils in this. *GCompris* has now been localised into more than 20 languages.

Space here doesn't permit a review of even a fraction of the promising Free Software apps that could have a place in our schools (see Debian's www.debian.org/devel/debian-jr/packaged and KDE's <http://edu.kde.org/> to see some of the best examples), but special mention must be made of two innovators: *OpenOffice.org* and Linux Terminal Server (LTSP).

WORLDWIDE COLLABORATION

OpenOffice.org (OOo) has been very active in involving school pupils in its community. The recent logo competition for the *OpenOffice.org* Schools Project (<http://marketing.openoffice.org/education/schools/>) had pupils from schools all over the world submitting entries. OOo also encourages school students to produce clip-art and other work for OOo, and sees this vital community involvement as part of encouraging the next generation of FLOSS users.

According to Ian Lynch, Schools Project lead at *OpenOffice.org*, "Andrea's contribution is not only beautiful and effective but also illustrates the potential of young people to contribute importantly to real projects under the banner of Free and Open Source Software. The idea is to learn by working together on projects such as *OpenOffice.org* whose license allows free collaboration across borders and between

schools. We invite all educators and students to consider involving themselves in OpenOffice.org."

If you've thought about participating, here are some primary aims that those interested in the project could consider:

- Join the QA (Quality Assurance) project and learn how to de-bug a world class software application.
- Contribute to the website by donating graphics and artwork
- Contribute to *OOoExtras* by providing clip art and templates – donate a brick and get a house!

It should be noted that – like a lot of Free Software – OOo also runs on Windows. In this guise, it can now be found in a lot of schools – and can be seen as a stepping-stone on the journey to eventually running OOo on GNU/Linux.

CHESHIRE PUPILS HELP OUT

Wilmslow High School got involved in supporting OOo overseas when it found out about the *OOEdu* project from postings by Ian Lynch in the uk.education newsgroups. On contacting Lynch, they were given the name of a contact, Colin Charles – the *OpenOffice.org* community marketing contact for Malaysia – to liaise with. A small group of year 11 students on a vocational ICT course were challenged to create and distribute 50 *OOo11* CDs for use by schools and students in countries where the cost of commercial software and limited Internet access would be a barrier to IT learning.

Initially, one group of students downloaded the OOo ISO from a mirror site. It was pointed out in the mailing lists that as the ISO didn't include *JRE*, it might handicap students with no Internet access, as some parts of OOo need it. Students downloaded *JRE* from the Sun website, but had problems adding this to a folder on the original ISO. Their solution was to extract the ISO, add a Java folder with *JRE*, and recreate their own ISO. They then successfully tested this on Windows and Linux machines before using it as a template for burning CDs. The discs are localised for all main Malaysian languages. Another group of students has created customised CD labels and wallets using the resources from the OOo website.

"From a teacher's point of view, the *OOEdu* project is ideal!" says ICT teacher Paul Burgess. "Our students get to learn about more than just IT! They're learning about citizenship, co-operation and working to tight deadlines. Because we've tried to make our project self-financing, they're also learning about the economics of production.

"The students have responded well to the project (even bringing in resources from home)," continues Burgess, "and are enjoying the responsibility they've been given. Knowing that they're making a meaningful contribution to an international project is a great motivator! For the future we've got plans to sell OOo CDs to parents and students at our school to subsidise the cost of creating more CDs for Malaysia. One of the students has suggested we try to get individuals and local organisations to sponsor individual discs, and this is something we'll definitely be considering."



Andrea Maggioni won the *OpenOffice.org* School Mascot Competition with this cheerful picture – community involvement in Free Software projects can be started in many ways.



From Malaysia, Charles adds: "What Wilmslow High School students are doing is excellent. So many schools have benefited from the exercise and they've even been kind enough to offer it to small businesses. OpenOffice.org is reaching more and more people daily, and WHS is helping that push; thank you, students!"

La Salle Secondary school in Malaysia is one of many schools able to stretch the IT budget further, thanks to the 00o CDs donated by Wilmslow High School in Cheshire – another example of 00o getting school pupils involved in the FLOSS community.

ARE YOU BEING SERVED?

Terminal services have long been a ready fit in multi-user 'nix environments, but are only just beginning to gain ground in the Windows world. The Linux Terminal Server Project makes the whole thing simple. All the applications run on one central server – easily backed up and protected from possible failures. Pupils' terminals contain no hard drive, only a boot ROM on their NIC, or a boot floppy. A kernel is fetched across the network, and remote-running apps are displayed locally on the (thin) client by the X server.

K12-LTSP is an adaptation of this for schools. Although it is an American project, there is a lot – particularly the case studies – that the UK education system could take from it. LTSP is advantageous for schools for a number of reasons:

- Old hardware (Pentium 90, 1MB graphics card) can be usefully and successfully deployed.
- Easy backups of pupils' work.



- Difficult for the students to break or mess up the setup – a simple power-cycle of the client, and you can log back in to a clean desktop.
- Money saved on clients can be reallocated for training and/or decent monitors.
- Easy to upgrade – just one machine to worry about.
- Scalable – just add another server and load-balance.

SCHOOL MANAGEMENT

Schools should have some flexibility about what software can be used in lessons. After all, the integration of computers into the curriculum is still at a relatively early stage, with many teachers ignoring them whenever they can, and little obvious benefit from them in some lessons. However, administering modern schools is complex task.

SchoolTool, from the Shuttleworth foundation, is looking very promising – but some work is still needed. Mark Shuttleworth has stepped in where governments have failed, and is offering bounties totalling US\$100,000 per year for desperately needed Free Software apps. *SchoolTool* reached another milestone mid-May 2004, with internationalisation and support for *Apache*-based virtual hosting. *SchoolTool* is coded in Python, and is very open to help in any form. Watch this one closely, the money and the will is there. Meanwhile, solutions are available for particular tasks – a search through the archives of the suse-linux-uk-schools mailing list will yield many apps. The venerable Linux for Schools Project (www.lfsp.org) has a niche, but comprehensive solutions must compete against big-budget rivals. This is one area where a little official encouragement for a FLOSS solution could make a great deal of difference.

A lot of work goes into planning lessons, and sharing resources can be a big help to teachers. *Moodle* is "a course management system (CMS) designed to help educators create quality online courses. Such e-learning systems are sometimes also called Learning Management Systems (LMS) or Virtual Learning Environments (VLE). One of the main advantages of Moodle over other systems is a strong grounding in social constructionist pedagogy." If you want to find out what that means, then feel free to visit the webpage. But what you really need to know is that teachers at home and abroad are readily adopting *Moodle*, as soon as they experience its power. *Moodle* is a 'killer app' for schools, and has impressed many who would otherwise have been unlikely to try Free Software.

EDUCATION EDUCATION EDUCATION

POUR ENCOURAGER LES AUTRES: A GEOGRAPHY LESSON

So, how do UK schools compare to those of other nations? Many countries have made headlines with commitments to FLOSS, but has it reached schools – and what of our nearest neighbours?

France has long had organised lobbying for Free Software, with APRIL (Association pour la Promotion et la Recherche en Informatique Libre) instrumental in five years of government backing for Free Software use. FLOSS support

is also widespread at the local government level.

Pingoo (www.pingoo.org) is a communication server distribution based on Debian, developed by the local government of Haute-Savoie. The ICT agency which



COVER FEATURE **EDUCATION**

deploys this technology, CRI74 (www.cri74.org) is an active FLOSS supporter. Pingoo is deployed at schools as a server for mail, file-sharing and user accounting, but the client-side remains MS Windows. CRI74 also organises LinuxEdu (www.linuxedu.org), and hosts OFSET (for more of which, see below).

SCEREN (www.sceren.fr) – aka CNDP network (<http://logiciels-libres-cndp.ac-versailles.fr/>) – is a national network of local agencies (named CRDP and CDDP) dedicated to pedagogical documentation and ICT in education. This task force, which is dedicated to writing documentation about Free Software, has mandated a free software company to produce a dedicated Debian distribution for the education sector – on the desktop side. Another CDRP project is SambaEdu (www.crdp.ac-caen.fr/forumse) – a Samba kit to ease set-up for schools.

Local communication server solutions abound, such as the Academie of Grenoble's Serveur de communications Linux pour l'Internet Scolaire (SLIS), which is deployed by over 600 schools, and EOL (<http://eole.orion.education.fr/>), from the government's national educational office. Of course,



Across the channel, MS doesn't get all the taxpayers' money. The LinuxEdu conference is organised annually every May by CRI74 – the IT agency of the local government of Haute-Savoie.

these are mostly server-side, but there is little doubt that the desktop is on the way and the political will is there.

France is also home to OFSET (The Organisation for Free Software in Education and Teaching) – an international group promoting and developing free software for schools. We have already mentioned its *GCompris* project; OFSET is also noted for FREEDUC – a bootable Linux distro with an educational bent. The last release (1.4) received UNESCO (United Nations Educational, Scientific and Cultural Organisation) funding to ensure internationalisation in English, French and Spanish, with printed documentation.

VIVA ESPAÑA

Of the 17 regions of Spain, 16 are planning to make significant FLOSS moves this year. Extremadura is a particularly good example, making news last year when its regional government sent out hundreds of thousands of distro discs to local residents, and started a wholesale switch-over to Free Software.

Andalucia is also pushing a major FLOSS initiative. Madrid even has its own distro: MAX – get it from www.educa.madrid.org/web/madrid_linux/archivos/obtener.html. Ismail Ali Gago, MAX Coordinator, believes the new release will “show teachers and students that GNU/Linux and Free Software programs are as easy-to-use and powerful as the others.”

Christian Einfeldt, who with director Paul Donahue made the film *The Digital Tipping Point*, available from the website www.digitaltippingpoint.com/, went to Extremadura and interviewed a number of those involved. One memorable comment in the film is “we now feel like they are ‘innovators’, whereas before the creation of Linex [Extremadura's Linux distro], we felt like there were merely ‘integrators’ of American technology.” Spanish ministers for education and for technology speak about gaining independence from US technology domination, and this is a powerful motivation for many countries.

In Italy, more than 1,000 Schools and Learning Centres are using OOo regularly; 263 of them even deliver courses and host Test sessions for the ECDL certification (European Computer Driving Licence) using OOo. On April 3 2004, Mario Pelosi – Head of the Italian Department for Innovation and Technology – appeared on behalf of the Italian Minister for Innovation and Technology, Dr. Lucio Stanca, to officially open the new PC Lab of Scuola Media Statale ‘G.Bovio’ of Foggia (www.smsbovio.com). This Lab is equipped with 18 refurbished PCs running K12-LTSP software. Italy's interest in LTSP follows a visit last year from Jim McQuillan, the lead developer of LTSP, to the First International EdOsNet.org Conference (www.EdOsNet.org).

The opening ceremony of the Lab was followed by a National Conference, held in the premises of Fiera di Foggia. It was attended by 230 teachers, coming from all over Italy. The Conference was entitled ‘Refurbished Hardware and Open Source Software for the Italian schools’. The keynote speech was delivered on video by Dr. Stanca, and was followed by the speech of Alessandro Musumeci, General Manager for Innovation and Technology of the Italian Ministry of Education. GNU/Linux desktops have arrived in Italy.

RESOURCES AND ACKNOWLEDGMENTS

LINUX PRO THANKS THE MANY DOZENS of individuals who have contributed information on FLOSS in UK and overseas schools for this article, and regrets that they cannot all be named individually.

Association For Free Software
www.affs.org.uk/education

Schoolforge UK www.schoolforge.org.uk

Schoolforge <http://schoolforge.net>

A not-for-profit project to provide a way of levelling up ICT skills
www.irlcomputers.com/ingots

Skegness Grammar School using Free Software in secondary education
www.scholarpack.org

British Educational Communications and Technology Agency www.becta.org.uk

North Chadderton High School: Bringing Educational Creativity to All (and cheekily pinching acronyms!)
www.becta.org/postnuke

Wilmslow High School's OOo project
www.wilmslowhigh.cheshire.sch.uk/ooo/

The SchoolTool Project
www.schooltool.org

Moodle's social constructionist pedagogy
<http://moodle.org/doc/?frame=philosophy.html>

GCompris
<http://ofset.sourceforge.net/gcompris/>

Ateb: providing and supporting IT solutions built using Free Software info@ateb.co.uk

Welsh localisation of 'nix
www.kyfieithu.co.uk

Powys' long-installed FLOSS solutions
www.telecentres.com/open-source/index.html

SCEREN's Debian-based schools desktop distro http://logiciels-libres-cndp.ac-versailles.fr/article.php3?id_article=18

Italian schools with OOo ECDL
www.aicanet.it/ecdl/OpenSource/Sedi_Open_Source.htm

Italian schools Open Source/hardware Conference www.progettolazzaro.it/ConvegnoFoggia.htm

SUSE Linux UK schools mailing list suse-linux-uk-schools-subscribe@suse.com

KDE edutainment project
<http://edu.kde.org/>

THE UK LAGS BEHIND AFRICA

Brazil is so whole-heartedly behind its move towards FLOSS that there is even a Minister for Free Software Conversion. India has been making convincing sounds about embracing FLOSS, and the Peruvian government response to Microsoft's overtures is well-known. In Africa, Namibia has more than half of its schools running GNU/Linux (featured on page 102, *Linux Format* issue 49). Sun Microsystems, meanwhile, is counting on sales of its Linux-based Java Desktop to China in the 10s of millions – so the UK and USA will soon be in a minority, fixated on teaching an expensive and outdated OS from an increasingly unpopular former monopolist.

Ian Lynch, of schools' supplier The Learning Machine, and education contact for both OOo and AFFS, agrees the point. *"I believe that the 'developed' world is very much behind 'less developed' nations in terms of FLOSS adoption, and I do believe that we are going to see a narrowing of the tech gap, and possibly even a reverse flow of tech development from India, China, Extremadura Spain, Brazil. If the 'developed' world doesn't get over its hang-ups with 'free' or 'libre' (or both), they are going to find out that other regions are eating its lunch."*

Lynch points out some of the benefits overseas schools are reaping. *"In Extremadura, there is one computer in each classroom for every two students! One kid does his/her work on the computer, the other does his/her work on paper, then they switch. That level of computer density is characteristic of only prosperous private schools in the US, or at least in California."* Lynch highlights concerns for the UK's IT industry, as it will face *"kids who are growing up in India, China, or Madrid"*.

Even in the USA, GNU/Linux has made inroads on the desktop of educational establishments – thanks to K12-LTSP. Paul Nelson of Riverdale High School in Portland, Oregon, is a tireless advocate. He saved the school some US\$50,000 when he installed the system on its 120 desktops, and is often found on the K12-LTSP mailing list helping others make the move. However, even Nelson had to install some MS software: *"We did pay \$10,000 (but only as a one-time cost) for MS terminal licenses for our Linux thin-clients, because we wanted our kids to have OS choice. Now, three months into the school year, they've made the choice and the MS stuff is not used much at all. I wish we could take that one back..."*

THE FUTURE

So, where does this leave us? We have many schools running FLOSS as part of what a recent *Linux Journal* SuitWatch article appreciatively called 'Do-It-Yourself IT' – traditionally a Linux strong point, but not good for schools lacking support, training and technical expertise. But what needs changing first?

"The major obstacle I see is the Microsoft annual licensing scheme", says Clive Menzies, Chair of Governors at Highgate Wood School and College of Performing Arts, an inner-London comprehensive school. *"The school lacks the resources or time to track software usage on every PC in the school, so it has opted to pay the Microsoft tax. Having committed to paying the annual fee, the incentive to look at anything other than Microsoft products is greatly reduced."*

The Office of Fair Trading (OFT) is currently investigating this licensing scheme, but it is problematic getting evidence



PHOTO: GARRY SADDINGTON

How long before the UK government stops pouring money down the drain, and lets our children learn that software isn't all about profit?

CONCLUSION

OUR CHILDREN – AND THE future of the once-mighty UK IT industry – are being let down by an education system that doesn't understand Linux, that hasn't grasped the educational benefits of software freedom, that hasn't realised an Open route to easy accessibility; inclusivity and community-building lies open before it. Hard-pressed decision-makers in the educational, business and government spheres with little IT expertise will continue to pick what they out-datedly believe to still be an 'industry standard'; unless they are told otherwise by us. Told simply. Told repeatedly. And again. Remember, it's *"education, education, education."* Go and make them listen.



"simply because the licensing is so complex and people don't necessarily understand it" says Ian Lynch. *"It's a bit ironic that the monopoly is so far-reaching that people don't realise there is an alternative, so finding evidence that they are not choosing the alternative they don't know about is difficult – when the monopoly means they don't know about it"*

Of course, even if the licensing scheme is disbanded, there remain further barriers to bringing Free and Open software into schools, to teach children the values of sharing and community over monopolistic greed and selfishness. Principally:

- 1** Political will – MPs need to be made aware of the issue, so that they can stiffen Becta's backbone.
- 2** Training – a generation of point-'n'-clickers need to learn a real OS, and stop being terrified of computers. Nobody likes change: technicians and teachers need coaxing through it. Also needs political will: see point 1.
- 3** Software – needs commissioning by government – we cannot be dependent on the Shuttleworth Foundation, a South African charity, to pick up the tab for our needs, while sending £100million+ per annum mostly overseas on software licences. Needs – guess what? – political will; again, see point 1
- 4** Education software suppliers – will not even fix programs to work with standards-compliant browsers such as Mozilla – in breach of e-gif guidelines. The political will to tackle this must be found: see point 1.
- 5** Lock-in – inertia is a powerful force. Lock-in can be prevented by enforcing demands for Open standards in contracts and software specs. Not surprisingly, this also needs political will: see point 1.

But let's not be too hard on Becta – we'll have to see if its recent rhetoric is matched by deeds in the coming year. Schoolforge-UK's Richard Rothwell assessed Becta thus: *"Becta may have had a bad press over the years with respect to FLOSS, but it seems currently to be awakening to the potential. With Prof David Hargreaves, Chair of Becta, speaking at the FLOSSIE Conference, and Becta's current project to evaluate the TCO of FLOSS solutions, we can detect a warning towards such ideas..."* ■■■

NO MONKEYING AROUND

PAUL HUDSON talks to two of the lead architects for Mono, Ximian, and SUSE, about Novell's future plans for Linux and Open Source...

Novell's acquisition of both Ximian and SUSE wasn't accidental or even co-incidental: it has a strategy for Linux that goes far beyond what most people realise. Novell's traditional strengths have always been in groupware and messaging, SUSE has both enterprise server products and desktop applications for consumers and workstations, and Ximian brings powerful desktop integration and management to the mix – if Novell is able to bring these all together so that the whole is greater than its parts, it should result in a powerful solution for businesses looking to Linux from top-to-bottom.

This is the first of a five-part series looking at each aspect of Novell's plans for Open Source, starting with how the acquisition of Ximian influences the strategy. Ximian produces Ximian Desktop, *Evolution* and Mono, but also employs several developers that are key to the GNOME development effort. As a result, the Ximian part of Novell is critical to the future of Linux on the desktop. We spoke Novell's Miguel de Icaza, Vice President of Development, and Richard Lindstedt, Linux Solutions Manager, about Mono and C#, Ximian Desktop, GNOME, and the big push towards the Linux desktop.



LINUX PRO: What kind of projects would you say could benefit most from switching to Mono?

MIGUEL DE ICAZA & RICHARD LINDSTEDT: To date, the Mono project has been very successful, with some Open Source websites suggesting that it will become the preferred platform for Linux development. Since Mono Beta 1 was released on 4 May 2004, more than 50,000 copies of the software have been downloaded.

Novell, certainly sees it as an important step in allowing enterprises to run a mixed Microsoft/Linux environment. In June 2004, we launched Mono 1.0, a joint community initiative to create an Open Source development platform based on the .NET framework. There are various kinds of Mono user, but they fall into two main groups: users who are creating new apps that are looking for a modern development platform, and developers who have an existing code base running on Windows and want to adopt Novell Linux Services.

Mono is an implementation of the ECMA and .NET platforms for Linux, Mac OS X, Unix, and Windows. A huge benefit of Mono is that you don't need to rewrite code, it can be gradually moved to a new platform by taking advantage of features, Mono embedding and platform invocation.

Mono expands beyond the original goals of the .NET Framework and gives developers access to various features unique to Linux such as tools to build desktop apps integrating with GNOME, the ability to re-use the *Mozilla* web browser, and tools to use OpenGL and its extensions.

LXP: How do you plan to tackle the move to Longhorn, and also newer releases of C# that incorporate improvements to the language, such as generics?

MDI & RL: The potential risk with Longhorn is that Microsoft could require those using APIs common to both Mono and Longhorn to pay license fees. While this decision rests with MS, what Novell is doing now is to provide implementations of Longhorn technologies such as Indigo and Avalon, new XML web services and user interface layers. Generics? We already have a version that we are shipping as a 'preview'.

LXP: Will parts of GNOME/Ximian be ported to C#?

MDI & RL: New applications developed within Novell for the Linux desktop are being developed with Mono and *Gtk#*, while old applications are using the mechanisms previously described. This allows for a nice migration that doesn't force a rewrite, and allows us to move at our own pace.

LXP: What about the desktop? Is there a roadmap for the continuation of Ximian Desktop (XD)?

MDI & RL: Novell Ximian Desktop is one of the most comprehensive enterprise Linux desktops on the market. Novell believes the Ximian and SUSE offerings are complementary, and that is exactly why Novell decided to acquire both companies last year. This was proved when, in January 2004, we announced that Ximian Desktop 2 had been updated to support SUSE Linux Desktop and SUSE Linux 9.0. The current version of Ximian Desktop 2 includes the Ximian Edition of *OpenOffice.org*, which reads and writes files in the latest *MS Office* formats, as well as the award-winning *Ximian Evolution* for email and personal information management that can be directly integrated with *MS Exchange* and other messaging and collaboration servers. The package also includes a *Mozilla*-based web browser, popular browser plugins, professional fonts, and built-in *ZENworks Linux Management* for software management.

In terms of areas for future development, Novell will continue to enhance the look-and-feel of the user interface, as well as integrating our Linux solutions with existing Novell back-end services such as *GroupWise* and *iFolder*. This will help significantly increase the functionality available for enterprise Linux, for example ensuring new software updates are made automatically, quickly distributing new applications company-wide or allowing users to access their information securely wherever they are.

MORE INFORMATION

Mono embedding:


www.gomono.com/docs/index.aspx?tlink=11@xhtml%3ahtml%2fen%2fembedding%2fphilosophy.html

Mono platform invocation:

www.gomono.com/docs/momodoc.ashx?tlink=11@xhtml%3ahtml%2fen%2fwrapping%2finterop.html

Beagle: www.GNOME.org/projects/beagle

Novell Ximian Desktop 2 is free for download at www.ximian.com/download



LXP: Are we likely to see it used as a standard in SUSE?

MDI & RL: As mentioned earlier, we believe the Ximian and SUSE solutions are complementary. Novell is currently working on our next-generation Linux desktop, which takes the best from SUSE Desktop and Ximian's and adds some new functionality from Novell to make a new product. Novell Linux Business Desktop is due to ship later in the summer.

You do have two different approaches in terms of GNOME and KDE, however. Novell doesn't see this as a problem: the most important thing is creating a desktop environment that's stable, cost-effective to administer and secure, and interoperable with existing network services. What we really want to do is break the cycle of proprietary desktop standards and offer businesses a genuine choice – whether they choose to use Linux throughout their business or maintain a mixed Windows/Linux environment.

Novell will continue to support both GNOME and KDE desktop shells, but we also feel that Open Source projects will benefit from closer partnership and collaboration to create a more unified set of desktop technologies.

LXP: What's the state of the Linux desktop market?

MDI & RL: Linux on the server side is now seen as a viable alternative to traditional operating systems. Big industry names like IBM, HP – and, of course, Novell – have resolved concerns about reliability, scalability, security, and ongoing support. Desktop Linux is the real one to watch. While Windows may dominate the market with a massive share, Linux is still the fastest-growing, second most-prevalent desktop operating system. The next year will be very exciting for those involved in the desktop space – Novell is introducing its Linux Business Desktop offering, and you'll see the first big global deployment of desktop Linux. Novell is moving 6,000 employees over to Linux. By 1 August 2004, everyone in the company is going to be on *OpenOffice.org*; and then, by this autumn,



GIVING BACK

What motivated you to open-source the Evolution Exchange Connector?

MDI & RL: Ximian was always known for focusing on usability and innovation in the ergonomics of the desktop experience, and Novell is committed to continuing this. The recent open-sourcing of the *Evolution Exchange Connector* demonstrated our commitment to openness, and was done as a means of furthering our efforts for interoperability with Windows environments.

For the IT department, technologies like *Connector* that allow Linux desktops to operate inside a mixed proprietary/Open Source environment, make choosing a Linux desktop a more reassuring choice for a business. With the release of the *Connector* source code, the entire *Evolution* product is now available under the GNU General Public Licence.

Open-sourcing the *Evolution Connector* allows *Microsoft Exchange Server 2000/2003* users to easily manage their email, calendars, schedules, address books, public folders and tasks from Linux desktops and was a vital step in reassuring everyday users that they can communicate, collaborate with their colleagues and customers.

roughly half the company will be on Linux, and the rest we are finishing off as soon as possible after that.

Clearly, there are areas that need addressing. The biggest consideration is about changing cultures and perceptions. Everyone is familiar with the Microsoft Windows and a lot of work has gone into creating a user friendly, easy to understand desktop Linux interface. However, business users may fear the unfamiliar: an education programme is needed to ensure people understand that different systems can be interoperable – you can transfer files from Windows to Linux; you can keep your email address book.

It's also important to provide backing to the developer community, so that the range of applications available on Linux continues to increase and it's possible to get bespoke desktop applications developed for Linux.

In terms of delivering the basic tools needed for business, desktop Linux is pretty much there. The only real limitation is for serious number-crunching in spreadsheets, but for the majority of day-to-day tasks, Linux is a realistic option.

LXP: Was *Dashboard* just a bit of research fun? The blog by Nat Friedman seems to have gone silent!

MDI & RL: The idea behind the *Dashboard* was to create a way for a computer to automatically show the user information and tools that will help them with the task they are performing. For example, while you read email, browse the web, write a document, or talk to your friends through instant messaging, *Dashboard* pro-actively searches for files or objects that are relevant to your current activity.

The *Dashboard* turned into a new project called *Beagle*: the first version was made available on 10 June 2004.

LXP: Big steps were made in usability (new file selection dialog, spatial *Nautilus* etc) and functionality (*Epiphany*). How do you plan to follow this up in GNOME 2.8?

MDI & RL: We're glad the recent developments in GNOME have been well received by the community. Novell sees GNOME as one of the most important projects in terms of enhancing the user experience of desktop Linux and allowing companies to run a mixed Windows/Linux environment.

Expect the next iteration of GNOME to further enhance the management functionality for the IT department and create an intuitive experience for the end-user. For enterprise users, we need ways to run Windows applications, support thin client devices and port existing applications to Linux and this is where you will see cross-over with Novell's existing resource management tools such as *ZENworks* and of course *Ximian ZENworks Linux Management*. ■■■



“WHILE WINDOWS MAY DOMINATE THE DESKTOP, LINUX IS STILL THE FASTEST-GROWING OPERATING SYSTEM.”

MIGUEL DE ICAZA, RICHARD LINDSTEDT

WIND RIVER WORKBENCH 2

NICK VEITCH greets the newcomers to embedded Linux.

While the strictly Linux-centric may not have heard of it, there are few bigger names in the world of embedded development than Wind River. A large selection of tools, software and services has made Wind River number one in many embedded markets.

But the recent announcement of *Wind River Workbench 2* (formerly known as *Wind Power IDE 2*) has also propelled the company into the Linux marketplace, by supporting development for Linux platforms, on Linux platforms, as well as allowing Linux to be used as a platform for developing VxWorks software.

Vxworks is Wind River's best-known product – a custom embedded real-time Operating System used to run many, many embedded devices. But though VxWorks has been phenomenally successful, the rise of Linux in the embedded marketplace simply can't be ignored. Embedded devices are one of those niches where there is a very good match with the benefits of using Linux – there are no licensing costs (important in the cost-conscious world of mass production), it's customisable (you can build just what you want into the kernel), it's scaleable, has a wide range of architecture support and has a mature set of development tools.

So what can Wind River offer with its *Workbench 2* IDE? *Linux Pro* spoke to Matthias Stumpf, Wind River's marketing manager for EMEA.

"Wind River Workbench 2 is an end-to-end solution: hardware bringup, firmware, application development, platform verification testing. This includes cross-platform development including multiple processors and multiple architecture. It supports any project regardless of the operating systems – Linux, different processors, cores architectures languages and target systems."

Stumpf adds: "Wind River Workbench optimises the tasks, especially the transition between phases of development. For embedded development, there are usually phases between hardware bringup and driver development to application development. Usually these phases use completely different tools and so managing these transitions causes problems.



From in-car infotainment to industrial devices, by way of consumer electronics – see Wind River's range of development tools on the website.

“WORKBENCH 2 IS AN END-TO-END SOLUTION – IT SUPPORTS ANY PROJECT REGARDLESS OF THE OPERATING SYSTEMS – LINUX, DIFFERENT PROCESSORS, CORES, ARCHITECTURES...”

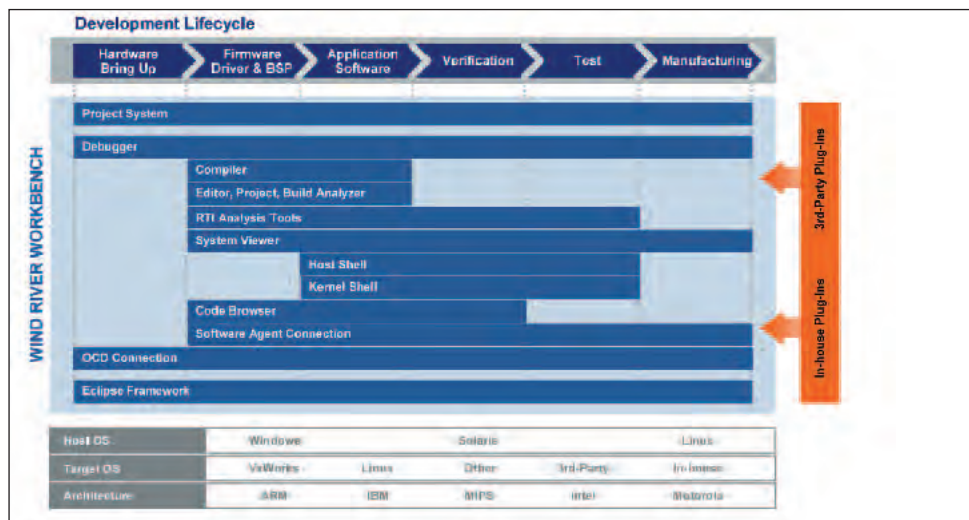
Later on, when the application developer runs into problems, he may think there is something wrong with the driver development, so that goes back and could go back to the board. Even if they all agree, you need to go back through the different phases."

ENVIRONMENT

Wind River Workbench 2 is completely based on the *Eclipse* environment (see box below for more details). Essentially the Workbench consists of a number of *Eclipse* plugins to provide editors, project management tools and a new debugger for the environment. The plugins provide functionality to cover different stages of development, from bringing up the hardware to writing device drivers and application development.

By using *Eclipse*, it is entirely possible to chop and change many of the tools used – it is possible to use the Wind River editor component, but compile on a standard *GCC* compiler for example, or use familiar *Eclipse* editors but target the Wind River compiler. This in itself can have benefits in terms of reduction in training costs, never mind that all development teams can use a single *Eclipse* platform no matter what projects they are working on.

The Wind River debugger is another key component as far as the company are concerned. Traditionally, Linux development has relied on the GDB (GNU Project Debugger) and KDB (Kernel Debugger) which compliment the *GCC* compiler. However, there are a number of practical problems working with these. When trying to trace multiple processes, some with many threads, many instances of the GDB are required which greatly increases system overhead, not to mention operator confusion. The new Wind River debugger is able to debug both kernel and system processes, enabling faster identification of problem areas of the implementation, as well as tracking all threads and processes. A special accelerated data retrieval system has been designed to help cope with the overhead of data transmission, which can be crippling on low-powered devices.



Wind River's idea of the typical development cycle.

The solution also offers a toolkit for analysing performance, including dynamic views of events, data, memory use and performance profiling. It also works asynchronously, allowing developers to continue working on the code.

A further plugin for on-chip debugging is expected to be included in the autumn release of the software. This will offer unique hardware debugging and testing for embedded devices. The features have been based on a study of Wind River's existing customer base and their work practices.

Matthias Stumpf went on to explain: "When we look at the development process, most companies seem to develop in teams using processes and tools that tend to be incompatible with each other. This tends to have an impact on cost, time and quality. The cost comes from having more than one tool – so there's licences, training and possibly increased headcount. This is one area."

"In terms of time, projects get delayed due to the issues of integration and troubleshooting the different areas."

While the use of separate specialist tools has been common practice, for embedded development they pose some real problems. Separate teams on separate systems have to work together, and very often problems can span the different layers of software being developed. But does all this integration really pay off for the customers?

Matthias Stumpf tells us: "We are seeing in some cases being able to reduce the development cycle by 20 or 25 percent, and we have seen that with some of our customers working with the Linux release."

SPECS

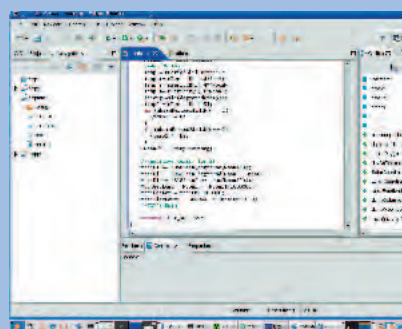
There are a few restrictions with the initial release. The first is that it has only been tested on Red Hat Enterprise Linux 3, so that is the only supported host platform, although there should be little difficulty in running it from other versions of Linux. The compiler supplied only currently supports the MPC82x series of embedded PPC processors from Motorola. This is a very popular series of processors, but is a little short of extensive – further architecture support is planned for future versions.

The embedded Linux space isn't exactly crowded, but with a massive rise in the use of Linux for everything from watches to smart network switches, it's a market that's definitely ripe for Wind River to carve a big slice from.



ECLIPSE IS A MULTI-PURPOSE IDE, BASED on Java. The project was started by a consortium which included the likes of SUSE, Red Hat, IBM and Borland, and more recently have been joined by MontaVista, SlikEdit and others. Although most popular as a Java development environment, it can theoretically be used for any language, thanks to a flexible plug-in architecture.

The *Eclipse* platform is released under the Open Source EPL licence, and is currently at release 3.0. The platform is a very extensible framework, which allows cross-platform development in a number of languages, with Java a particular favourite (the *JBoss* IDE is based on *Eclipse*). Functionality is added by a huge (and growing) range of free and



commercial plugins, providing a diverse range of tools for more-or-less every conceivable task: from spellcheckers and syntax analysers to complete language environments.

You can find out more about the *Eclipse* project at www.eclipse.org