

LXF53
Linux Format May 2004
● ULTIMATE LINUX BOX ● MANDRAKE 10.0 ● MONO ● UNREAL TOURNAMENT 2004 ● C/C++ IDE ROUNDUP ● KINO TRICKS ● SUBVERSION ● PLANMAKER ● GIMP2

LINUX FORMAT

THE UK'S BEST-SELLING LINUX MAGAZINE!

MANDRAKE 10.0
REVIEWED INSIDE
AND ON YOUR
COVERDISCS!

BUILD THE ULTIMATE LINUX BOX! COMPLETE 12-PAGE CONSTRUCTION GUIDE



DO MORE WITH LINUX!
Using Ximian Evolution **p62**
Network admin with PHP **p72**
Video editing with Kino **p70**
Using SQL databases **p76**

UNREAL GAMING
The biggest game to hit Linux since, er, UT2K3. Charge your weapons! **p22**

EXPLORING SNMP
Administering SNMP devices with Free software **p84**

C/C++ IDE ROUNDUP
The very latest development environments compared **p32**

MONO: .NET FOR LINUX
The Open Source option for .NET platforms explained **p56**

"Sun's commitment to open-source Java would speed the development of a first-class Java implementation"
Rod Smith, IBM's VP of Emerging Technology wants Sun to join the party... **p07**

DVD issue also available Printed in the UK
LXF53 MAY 2004 £5.99
future publishing
Media with passion 9 771470 423019 05
(£5.99 outside UK and Republic of Ireland)

www.linuxformat.co.uk

What's important?

With *Linux Pro* and *Linux Format*, we are constantly in touch with readers from all spheres of the Linux world.

Developers, sysadmins, desktop users – there is a rich and varied world inhabited by people using Linux. But it's a mistake to think that everyone sees things the same way...

For many enthusiasts and desktop users, the choice of Linux can be a purely financial one, or a decision motivated more by ideology: 'free as in beer, or 'Free' as in speech. Though the common factor is that they are using Linux, the reasons are different, and so are their expectations and views on many subjects relating to the technology. Some will steer clear of MP3 files in favour of OGG because of the patent issues surrounding the former; others may make the same choice because they reckon the latter gives better quality.

In the world of commerce, things become more interesting. When it comes to device drivers for example, as we saw recently in our *Get All Your Hardware Working* issue (LXF51), it is often a

simple question of economics for companies faced with the question of whether to support Linux or not. They aren't interested in ideology, they are interested in money in the bank (and why not?), so the question of Linux support isn't about 'doing the right thing', it's just about how much it would cost against the return. In some ways, with the Open Source community doing such a great job, it doesn't matter that much to them. And some companies that DO produce Linux drivers – but do so in a proprietary way – will draw flak from some parts of the community.

There certainly exists a pragmatic element to the Linux community, who want to use solutions that work, whatever licence they use or wherever they come from. But – whichever type of user you are – it's certainly important to remember that there is a Linux community, however disparate some of its parts may be, and to try and understand each other. That's what we try to do at *Linux Format* – you're the best judges of whether we actually achieve it.



Nick Veitch EDITOR

What would your ultimate desktop consist of? Find out here! p44

The first 2.6 distro is here, reviewed and rated and included on the coverdiscs. p18

Unreal Tournament 2004 – the best Linux game since Unreal Tournament 2003, probably p22



18



22



44

LINUX
FORMAT

AIMS OF THE MAGAZINE

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

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

MEET LINUX FORMAT'S TEAM OF WRITERS...



Andrew Channelle
LXF's newshound and the Beginners' best friend, Andy organises his entire life with *Evolution* this month.



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



Mike Saunders
Spends so much time scouring the Net for software for *Hot Picks*, we wonder whether his legs are still functional...



Jono Bacon
A core KDE developer, web developer, writer, musician and sound engineer, Jono also knows IDEs inside-out.



Paul Hudson
As if being the master of PHP wasn't enough, he wants to be the company *Unreal Tournament* champion too!

Maurice Kelly
Busy coder, electronic engineer and *Midnight Oil* fan Maurice reckons he's found the perfect replacement for CVS.

Tom Wilkinson
Who better than this long-time professional Linux expert to pick at Mandrake 10 until all its scabs fall off?

Chris Brown
When it comes to packets and protocols, Dr Chris knows what's what, and isn't afraid to draw a diagram.

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

Roger Burton West
Holstering his NERF gun, he asks whether yet another paid-for proprietary spreadsheet is a wasted effort.

CONTACT US

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

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

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

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

General enquiries:
linuxformat@futurenet.co.uk

Website: www.linuxformat.co.uk

Or send your letters to:
LINUX Format, Future Publishing,
30 Monmouth Street, Bath, BA1 2BW
Phone: 01225 442244
Fax: 01225 732295

More contact info on p114

Newsdesk

- Macromedia Flash for Linux ● HP Linux hardware ● GTK+ update ● Fedora drops XFree86 ● RSS
- Small motherboards ● Motorola Linux mobile phone ● Training Camp course ● Latest SCO News

COMMERCIAL LINUX DISTRIBUTION

Novell delivers its first SUSE upgrade

Novell/SUSE made use of Germany's enormous CEBIT show in Hanover to launch what it claims is the world's first complete commercial Linux distro based on the latest kernel. SUSE 9.1 will be available early May 2004 in both Personal and Professional iterations on 32-bit and 64-bit systems. In addition to a 2.6 series kernel, the distribution also features the latest official releases of KDE (3.2.1) and GNOME (2.4.2) and an expanded range of commercial, Open Source and demo software.

Novell acquired SUSE earlier this year for an estimated \$210 million in cash in what CEO Jack Messman called a "strategic shift" for the company. The buyout, together with the earlier purchase of desktop developers Ximian, gives Novell a complete end-to-end Linux solution and should provide solid footing for SUSE's entry into the US market, which has been dominated by Red Hat. It also secured a \$50 million investment from IBM and the ever-lasting ire of SCO.

Markus Rex, Novell/SUSE's General Manager, said the release would build on the project's reputation for power and ease of use. "Seasoned Linux users get the power of 64-bit computing and more than 2,500 software packages," he said, while beginners can take advantage of the Live CD included with 9.1 Personal, which allows users to test



Richard Seibt, Jack Messman and Chris Stone. SUSE 9.1 will be the first release under the Novell banner, and the "first commercial 2.6 product".

the system without affecting their PCs before committing to a full install.

In addition to the usual application suite – covering all the bases such as web browsing, email *etc* – SUSE 9.1 also includes for the first-time demos of *TextMaker* and *PlanMaker*, a word processor and spreadsheet app respectively, from SoftMaker, as well as *Rekall* for database access, a personal finance manager called *Moneyplex* and the latest stable version of the *Scribus* desktop publishing application.

SUSE 9.1 Professional adds server functionality to the equation and comes on five CDs and two double-sided DVDs to cover 32- and 64-bit installations.

64-bit users will find systems tailored for both AMD64 and Intel Extended Memory 64 Technology. The Professional package will also include for the first time *Samba 3*, which allows seamless integration into Windows domains with support for Active Directory.

Finally, with the view to establishing its Open Source credentials once and for all, Novell has announced a plan to release SUSE's *YaST* set up and configuration system under the GPL. The intention, sources said, is to allow companies such as Computer Associates and IBM to integrate *YaST* into their own offerings and, hopefully, see it adopted as an industry standard



Mandrake 10.0 will form the basis for the next version of Mandrake Corporate Server.

way for managing software and hardware on Linux systems. It is also hoped the move will attract external developers who will add support for rival distributions such as Fedora and Debian, giving the *YaST* module more traction in the wider community.

Since its debut, SUSE has resisted all requests to Open Source the tool, citing the commercial advantage it provided in a crowded marketplace. Novell's move, in tandem with its efforts against SCO, will not harm the company's profile as a genuine Open Source player.

SUSE Personal, which includes two CDs, an installation guide and 30 days support, will cost US\$29.95, while the Pro edition, including administration guide and 90 days of support will cost US\$89.95. Also, a Professional upgrade version will be available for US\$59.95.

The French connection

SUSE's claim to the first 'commercial' 2.6 release may be tested by MandrakeSoft, emerging from the shadow of bankruptcy protection, with the launch of Mandrake 10 Official which, if all goes to plan, should also be

OPEN SOURCE Unified Java would challenge .NET

IBM has begun a campaign to convince Sun of the merits of bringing its Java technology to the Open Source community. Unusually – and perhaps fatally – IBM kicked off the initiative with an open letter to Sun from Rod Smith, Vice President of Emerging Technologies. The letter was a response to a comment from Sun's Simon Phipps asking why IBM hadn't Open Sourced its own implementation of Java. Smith took this, he wrote, as a suggestion that the two companies could work together on an open version of the technology.

"IBM would like to work with Sun on an independent project to open-source Java. Sun's strong commitment to open source Java would speed the development of a first class and compatible open source Java implementation to the benefit of our customers and the industry," Smith wrote. *"IBM is ready to provide technical resources and code for the Open Source Java implementation, while Sun provides the Open Source community with Sun materials, including Java specifications, tests and code."*

Sun's stewardship of Java has come under question recently amid claims that its firm hand on the technology is stifling its ability to compete with Microsoft's .NET framework. Studies suggest three million high-end developers are actively involved with Java, while .NET lays claim to eight million. One report by Precursor said offering the technology on an Open Source licence would remove a significant barrier to adoption, but that Sun's reluctance to be 'more than just a box shifter' would make it hard for the company to let go of the reins.



available early May. After the CD-frying problems experienced with Mandrake 9.2, the company made changes to its development roadmap, with 'Community' editions of the distribution coming out around two months before the 'Official' boxed versions (see *News*, last month). Additionally, the Official release will be followed by an update to Mandrake's Corporate Server product.

The emergence from bankruptcy protection means MandrakeSoft stock can once again be traded on the European Stock Exchange. The company will implement a debt reorganisation programme which entails the repayment of debt amounting to some 4.1 million Euros (without interest) over a nine-year period.

GENTOO'S QUARTERLY OFFERING

Hot on the heels of the year's first release (2004.0), the Gentoo project has scheduled a maintenance update for late April 2004. As before, the distro will be available on multiple architectures including x86, AMD64, PowerPC, Sun SPARC and SGI MIPS. The last release also saw the addition of a 'security-hardened' x86 implementation.

The next release is expected to include kernel 2.6.3, KDE 3.2.1 and GNOME 2.4.2, and the innovative Catalyst 'meta-tool' that allows developers and users to customise every aspect of their Linux system, from the creation of Live CDs, customised install routines and bespoke binary packages.

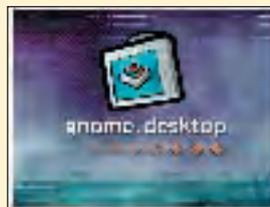
2004 also saw the Gentoo project enter a more organised phase with an online store, where users can donate to the project using a number of options including PayPal and Visa. Also available is a Gentoo subscription, offering each quarter's release delivered on CD straight to your door. <http://store.gentoo.org>

NEWSBYTES

■ **KDE 3.2.1** is the first maintenance release of the latest desktop environment from KDE. As well as improved localisation – KDE is now available in 49 languages including Bengali, Low Saxon and Latin Sebian – the new release also include fixes for a number of bugs. Source code and binaries are available from the usual sources.

■ **Vignette** has become the latest ISV to bring their software range to Linux. The company has added to its Linux roster by porting both *Vignette Content Management* and *Vignette Business Integration Studio*. Previously only *Vignette Application Portal* was supported natively. The software is certified for use on Red Hat Enterprise Server. Spokesman Mike Evans said customer demand had underpinned the company's choice. *"Organisations are relying on Red Hat Enterprise Linux for mission-critical applications that drive business efficiency across the enterprise,"* he said.

■ **PeopleSoft** is continuing its sterling efforts to bring applications to Linux. The latest addition to an already sizeable roster is *EnterpriseOne* which covers areas such as supply chain, supplier relationship, financial, asset life cycle and project management. *EnterpriseOne* version 8.10 is scheduled for release in the second quarter of 2004, and is certified for Red Hat.



■ As *LXF* went to press, the first release candidate of **GNOME 2.6** (codenamed *We Are the Champions*) was hitting the mirrors. The full version – which may or may not be different to the RC – should be available as you read this.

■ **IBM** has pledged to migrate a significant portion of its internal desktop systems to Linux by the end of 2005. Current figures suggest that 15,000 of the company's estimated 300,000 internal computer systems are running Linux, with plans afoot, according to International Data Group (IDG) to increase that number to 40,000 by the end of the year.

■ Michael Robertson would like your assistance in coming up with a new, temporary name for **LindowsOS** to be used in countries where Microsoft has received an injunction against the company. Send suggestions to temporaryName@lindows.com

Jono Bacon

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



COMMENT

Pushing the community

“ We are all well aware that Linux is a community-based system. This typically Internet-based community toils away to advocate the software that we use. This is great, but does not seek the true potential of Linux advocacy.

What I would like to see is a true local community of Free software and Open Source awareness. Though there are many people advocating Linux on the Net, this advocacy is only practical if you know where to look. Linux advocacy on Slashdot and Linux Today is pretty redundant: the choir already sing the same song. Advocacy needs to be a practical, community-based effort to get local community members in front of a Linux machine to see it and explore it. Getting a regular joe to see and use Linux is ten times more useful than writing about how great Linux is on the Net.

Don't get me wrong, I am included among the guilty here. I write about advocacy every month in *LXF*, but I hope these words of advice help you to push your promotion of Linux in your local community. I know that LUGs across the world try to push Linux in the face of managers, spouses, friends and others, but what about those not fortunate enough to know a Linux user? These regular folks whose computer experiences are typically crippled with misbehaving Windows boxes are lost outside this sea of advocacy.

Let us remember that Linux is free. Make use of this fact by getting CDRs of Knoppix into people's hands; putting people in front of a computer at PC World that is booted with Knoppix is an instance of advocacy with potential for *real* impact... ”

FLASH ON LINUX

Macromedia testing the water for Linux

Macromedia, developer of some of the most widely used *de facto* standard web design applications, is said to be readying a native port of its flagship *Flash MX* product. The company's chief software architect Kevin Lynch said during his keynote speech at the Flashforward 2004 conference that the recent addition of Flash support to *CrossOver Office* from CodeWeavers, undertaken with assistance from Macromedia, was being watched with interest by company insiders, and that future versions of the application would ship with *Wine*-specific optimisations.

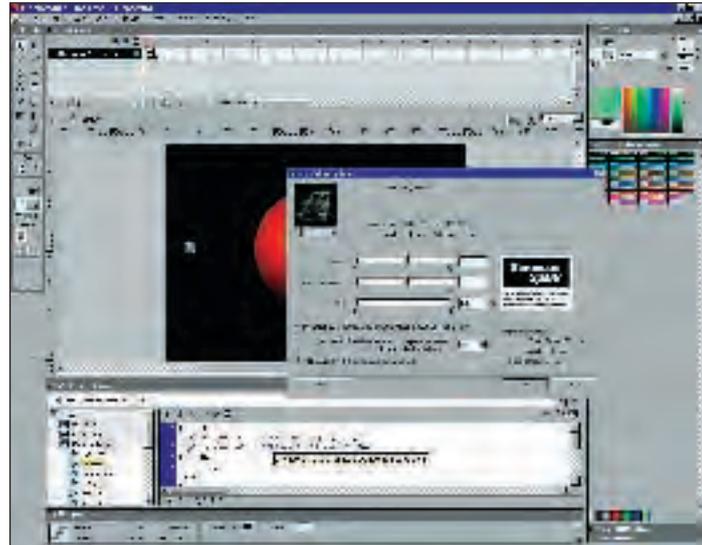
If this pilot is successful, Lynch said, the next step would be to produce a native version which could then lead to broad support for the

rest of the company's popular application range.

"What we've been investigating is, when it will be time to bring our tools to Linux," he enthused. "I think it might be now."

Macromedia is known to be keen to establish *Flash* more firmly in the marketplace – especially outside the design fraternity – in readiness for more competition from the emerging SVG standard; and with Linux approaching or surpassing (depending on which survey you believe) Mac OS as a desktop operating system, market pressures may see more 'household names' ported.

Macromedia's current Linux lineup consists of server products such as *ColdFusion* and the *Flash* plugin for *Mozilla* and *Konqueror*.



Flash already runs under Linux with CodeWeavers' *CrossOver Office*. How long before a native Linux version hits the stores?

Linux Web Watch/

It's my news!

RSS is a technology whose time has come, and inspired by Ximian Evolution's basic RSS options, LXF went in search of some more fully-featured news aggregators.

RSS readers falls generally into two camps, they are either standalone apps or use any available web browser to display content. The former tends to



have more features and is a touch faster in use, while the latter can be accessed on any connected PC, regardless of processor, OS or location.

Perhaps the finest RSS app for Linux is *Straw* (www.nongnu.org/straw/), which is part of the GNOME project. The latest release, 0.22 has refined the three-paneled UI from previous releases, but still has a stack of dependencies.

However, if you have an up-to-date GNOME install and Python, most of them should be met as standard.

KDE users are well catered for with *KNewsticker*, which is accessible on the *KPanel*. This adds a

« *Straw's* dependency problems don't obscure the potential of a fine app in the making...

scrolling news ticker on the panel, with each headline linking to the site (www.student.uni-oldenburg.de/frerich.raabe/knewsticker/). Furthermore, the ticker also has a menu which dynamically updates to reflect changes in your news sources.

One of the best web-based RSS clients is *Bloglines* (www.bloglines.com) which requires registration before you can begin assembling your sources. We noticed a slowdown when monitoring a large number of news sources, but this is negated by the service's accessibility. A hybrid of the above two methods is the *Mozilla*-based *Aggreg8* (www.aggreg8.net), a university project by Phil Roche. This is a nice-



» *Bloglines* features the amusingly named 'blogroll' in a totally chortle-free manner.

looking applet (if you like the *Mozilla* aesthetic), but we had trouble getting it to install on the latest *Firefox* release. It works OK with the full *Mozilla* suite though, and can be used on its own, or through a *Mozilla* browser window.

NEWSBYTES

■ **GTK+**, which underpins the GNOME project, has also received a recent update. GTK+ 2.4 includes a number of significant improvements including a 'radically simplified' file selector designed to make day-to-day use more intuitive.

■ **Hyundai** has followed Volvo's lead and installed a Linux-based IBM system to model vehicle-crash data, using a cluster of smaller computers to add super computer-like performance to its design process.

■ **VueScan**, the high quality, cross-platform scanning app is now available free-of-charge to download for personal/ educational use at www.hamrick.com. The package, supports a wide range of desktop and professional-quality scanners, including many not supported by Open Source software. A commercial licence for the standard edition usually costs \$59.95, while the professional version retails at \$79.95 per user.

■ Intel has finally, though not yet officially, made Linux drivers available for its mini PCI wireless adaptor. The IntelPRO/Wireless 2100 device is best known for its inclusion in Centrino notebooks. See <http://ipw2100.sourceforge.net/> for more details.

■ MapuSoft has created an OS abstractor which allows developers working on embedded apps to test apps across a range of Linux flavours. Raj Johnson, President and CEO, said the time was right to bring the product to market, as Linux was poised to become a significant element of the embedded market. "OS Abstractor gives developers a common interface to quickly start developing applications without having to worry about the OS variations across various Linux distributions," he said.

■ **Tsunami Research** has launched what it claims is a new spin on the clustering idea. Hive technology, named in the style of *Star Trek's* Borg race, adds a novel temporal layer over the normal clustering technologies. This makes it especially suited to financial tasks as each node in the cluster is aware of the time and can intervene in specific processes, adding more power, if time-sensitive processes are getting bogged down in the system. The Hive also includes a selection of self-healing tools and an abstraction layer designed to make the system more intuitive for programmers familiar with C++ and Java.

■ The efficiency of the **Linux kernel** development process has apparently increased two-fold since Linus Torvalds controversially adopted the Bitkeeper code management system in 2002. Red Hat's David Miller says the system allows him to spend more time working than wallowing in 'merge hell'.



LINUX HARDWARE

HP to offer Mandrake and TurboLinux

After months of dithering, Hewlett Packard has finally decided to offer – and more importantly promote – Linux as an option on a range of desktop computers in both the US and Asian markets. The PCs, aimed at the small and medium enterprise (SME) market, will come with either Mandrake Linux (US) or TurboLinux (Asia) installed.

The US initiative includes two models, the basic Compaq Business Desktop DX2000 features either a Celeron or Hyper-threaded Pentium 4, while the DC5000 will ship with HP's own Lifecycle management tools. Mandrake will be offered as an option on both platforms, though at the time of writing, all the pages on HP's US website with details of Linux machines rejoiced under the heading "HP recommends Microsoft Windows XP Professional". It seems that HP maybe still has a way to go in developing its Linux promotion strategy

beyond a sentence stating "Microsoft or Mandrake Linux promotes workplace versatility..". The range will be joined in the summer by the 7000 edition which will include additional security apps.

Meanwhile, in the Far East and India, plans have been laid to begin shipping TurboLinux PCs within the next few months, though HP is said to be concentrating on far bigger fish hoping to migrate thousands of seats to Linux. Some analysts have predicted that a more belligerent attitude from HP, the world's biggest supplier of desktop systems, could have a bigger impact on the popularity of Linux on the desktop than IBM's recent moves. Roger Kay, Vice President of IDC, said there seemed to be a momentum building up. "When you've got HP and IBM championing Linux on desktops, I'd say the fighting gloves have come off, and it's going to be a great battle."



HP's latest business desktops offer Mandrake 9.2 as a pre-install option.

Hoyt Duff

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



COMMENT

If Wishes were Horses...

“...Beggars would ride. Aroused by the onslaught of anxiety produced by the promiscuous permeation of the plethoric corporate world into Linux, I engaged in introspection on interesting innovations that Linux has introduced and wished about how they might be employed and enjoyed.

The CD-based distro is compelling, as these clever creatures permit some manner of Linux system to be launched and run solely from a CD-ROM. Noted among the newest is Knoppix, now spawning many variations suited for general desktop use as well as those that specifically focus on firewalling, forensics, security, and multi-media. See <http://lwn.net/Distributions/index.php3#cd> for an expanded list.

The first of two critical components are the undergirding technologies that permit file compression, loopback mounting and such that make it possible to pack a lot into a small package. The second is the sophisticated scripting and autodetection that cause the final result to be almost magically displayed before your eyes. It is astounding as to how advanced the technology has become.

While alluring, however, one needs a fast CPU, fast CD-ROM drive and lot of fast memory to make many of these distros useable for more than a stunning demonstration. What I wished for, however, was a CD-based distro that would exhaustively catalog the computer hardware, benchmark it, and demonstrate its compatibility with Linux. How handy this would be to evaluate any hardware, but especially laptops. Oh, joy! I would ride in style.

AGAINST THE ADVERTISING CLAUSE

Fedora puts forward plan to drop XFree86

After last month's ructions over licence changes introduced in the latest version of Xfree86, the Red Hat-backed Fedora project appears to be going one step further than Mandrake by including a rival X implementation in the next Fedora Core release. A number of distribution vendors, including Mandrake and Debian, last month said they would roll back to Xfree86 4.3 in order to avoid the advertising clause included in version 4.4. Fedora developers have taken the issue forward, apparently preparing RPMs for a Fedora Core 2 based on the rival/forked X.org project which is hosted by freedesktop.org. X.org was

forked just before the contentious licence adjustment (4.4.0-RC2), so compatibility with the current range of drivers is maintained. However, the move could mean more widespread changes in terms of support for the original XFree86, with graphics card manufacturers following the lead of the biggest distribution packagers to support the forked X implementation.

While the current dispute has been characterised as a licensing issue, speculation continues that the core developers of XFree86 and distribution vendors have been pulling in different directions for some time and the advertising clause was just a convenient method – perhaps even for both sides – to assert their position.



Fedora's decision may impact on the entire X development effort.

Embedded Linux News



● **Motorola** is set to launch its first Linux-based mobile phone in the US. The Motorola E680 supports MP3 music, MPEG 4 video, RealPlayer multimedia and MIDI ringtones, Bluetooth, GPRS and USB connectivity and SD cards up to 1 GB, making video on the move a real possibility. Motorola's previous phones built around Linux and Java have been marketed only in Japan and China, while more established software featured in US and European devices. Spokesman Geoffrey Frost said, "The device formally known as the cellphone is now our wireless entertainment portal." It

also features a camera with 8x zoom capable of capturing 640x480 stills and MPEG4 video.

● The installed base of **Linux-based set top boxes** is expected to top 40 million units by 2008 according to recent survey by Strategy Analytics.

● **LynuxWorks** is claiming to be first to market with a 2.6-based Board Support Package for Apple's G5 desktop. The company says adopting the Apple route should allow embedded developers to save money over IBM's own PowerPC 970 development boards.

● **General Micro Systems** has shoe-horned a complete hardware/software solution into a space 30 per cent smaller than a standard credit card. The Spider is available in two configs: a low-powered version based on IBM's 400MHz 440GP, and a more rounded version built around an 800MHz PowerPC 440GX. The boards are capable of booting either VxWorks-Tornado II or Linux directly from the onboard flash memory.



MINI-ITX

VIA gets even smaller with Nano

Already big in the world of small, VIA has released a new mini-sized motherboard called the Nano. Nano-ITX boards cover a mere 12x12cm, ideal for cramming into ever smaller boxes.

The Mini-ITX range of motherboards have already proved remarkably popular with Linux self-builders, who have manufactured them into all sorts of innovative and useful designs. With low-power requirements, but essentially all the onboard connections you would expect from a modern PC, the boards such as the EPIA M, are ideal for making home media devices such as PVRs, tiny file servers or firewalls.

The initial Nano board includes the new CN400 on-board graphics chip which includes hardware MPEG4 and MPEG2 decoding. With support for up



to 1GB of fast DDR400 memory, onboard RNG, SATA, USB, LAN and six-channel audio, the Nano should be just as big a hit with tiny systems-builders PC-ing up their toasters.

Get the full specification of the new Nano-ITX mini motherboard from VIA's embedded website at www.viaembedded.com/product/epia_N_spec.jsp



By Pamela Jones, www.groklaw.net/

■ What a month of activity this has been in SCOville. There were almost two-dozen major events in the last month alone. Both sides are building up steam in the legal battle, but SCO is losing the FUD battle.

■ Eric Raymond published a memo from Mike Anderer revealing that Microsoft arranged for the BayStar funding that kept SCO afloat financially and provided the war chest funding SCO's 'end user' lawsuits, two of which were filed this month against AutoZone and DaimlerChrysler. Anderer wrote: "I realise the last negotiations are not as much fun, but Microsoft will have brought (sic) in \$86 million for us including Baystar. The next deal we should be able to get from \$16-20, but it will be brutal (sic) as it is for go to market (sic) work and some licences. I know we can do this if everyone stays on board and still wants to do a deal!"

■ SCO PR man Blake Stowall acknowledged the memo as being authentic, but claimed it didn't mean what it seems to mean, as Anderer misunderstood the situation. However, BayStar later admitted that senior Microsoft executives had indeed asked them if they'd care to invest in SCO.

■ All is not what it seems on the end-user lawsuits either. If you read the fine print, neither lawsuit is exactly a 'Linux end-user' lawsuit as was promised in public SCO statements and in the Feb 11 S3 filed with the SEC. *MS Word* metadata in the DaimlerChrysler complaint indicated that originally SCO intended to sue Bank of America in a true end-user copyright infringement under the DMCA. At the last minute, that lawsuit was dropped, DC was plugged in as the defendant and the offence was changed to a contractual one.

■ AutoZone's crime, according to the complaint, was using SCO's libraries when it switched from Unix to Linux.

However, the man who did the switch left a comment on Groklaw shortly before the lawsuit was made public, claiming he didn't use SCO's libraries. It seems to be 'SCO customers' rather than 'Linux users' that need to be wary at the moment.

■ SCO also announced a new sign-up for its IP license: EV1server.net, an ISP that also hosts Linux and Windows servers. It was soon revealed that EV1 featured on MS's 'Get Linux' website last September extolling the wonders of Windows. All this did little to distract from SCO's quarterly results: only US\$20,000 from SCOSource initiatives this past quarter, yet it spent millions in the effort.

■ Finally, software author Fyodor revoked SCO's right to distribute Nmap. Because SCO has violated the GPL, Fyodor argues, terminating their right to distribute the security application is within his rights as its author.

■ For a full digest of this month's events, see www.linuxformat.co.uk

David Cartwright

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



COMMENT e-smith

“ I came across an interesting new Linux beast this month. It's called e-smith, and it takes an interesting side-step away from standard Linux distros by providing the functionality of a general-purpose mid-range server without the management trickiness.

Let me explain. With a traditional Linux distro, you're provided with a variety of tools (either command-line or GUI-based ones, depending on your flavour and whether you installed X-Windows) that allow you to manage the various applications. e-smith is different, as although you can get adventurous and dig around on the command-line, the majority of the stuff that the average user of this type of system are handled via a single Web GUI. So creating users, sharing disk areas, setting up websites, handling email, even advanced stuff such as integrating with an LDAP directory service or making the unit's *Samba* installation into a domain controller is all handled via this single, intuitive Web GUI.

A variant of the package has been released commercially by Mitel – the difference being that e-smith is freely available, albeit with no formal tech support (we're used to that with Linux!). It's idiot-proof to set up, and that got me thinking: "Maybe this is another route into the general OS market for Linux." After all, one of the things that puts people off Linux is that even the installers can be a bit scary, and then once it's set up you have to muck about with config files to make stuff like *Samba* work.

So if you're thinking of running up yet another Red Hat box to handle your mail and fileserving: give e-smith (www.e-smith.org/) a bash. Linux without the configuration hassle can't be a bad thing!

SECURITY TRAINING

Thinking like a hacker...

Training Camp, the accelerated IT learning company we covered in August 2003's *Linux Pro* (with LXF43), has since built upon its LPI training course by adding a new 'ethical hacking course'. After spending five days studying the ins and outs of hackerdom, successful students earn themselves the title of 'Certified Ethical Hacker', and will hopefully be much better equipped to protect their businesses against hackers. Robert Chapter, the co-founder of The Training Camp, said, "the best way for organisations to prevent cyber-damage is to learn the tricks of the real experts – the hackers themselves. The Certified Ethical Hacker course will give security professionals an insight into network vulnerabilities from the attacker's point of view."

As it has so much Linux experience from already teaching the LPI course, Training Camp is also offering a special 12-day course that encapsulates all the LPI and all the CEH tutorials and examinations in one sitting, thus



Training Camp is based in Tadley, between Basingstoke and Reading.

minimising time spent away from the office. The new CEH course is similar to the LPI course in that students stay at The Training Camp's dedicated student housing, as well as having all their meals and snacks provided – such centralised organisation helps students concentrate on the task in hand.

As hackers (or perhaps more pedantically – crackers) often use Linux or a BSD as their jump-off point to breaking into other servers, the combination of the LPI and CEH courses seems logical and is likely to prove popular – we're hoping to bring you a special *Linux Pro* update soon.

Mailserver

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

★ Letter of the month

This month's winner receives a copy of *Applied C++*

KDevelop

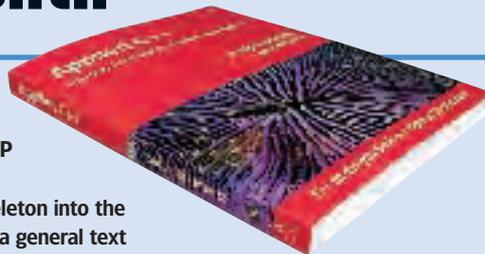
I found your tutorials for *Kylix* quite interesting, though being more interested in C/C++ than Pascal, I didn't follow them much beyond the first few.

Several times now, you have selected *KDevelop* to be among the top IDEs for Linux, so I think it would be a good idea to start a tutorial on building KDE-based apps with it. I know that it provides enough of a skeleton to compile a 'working' app to build

on. From this, you could touch on the basics of OOP design and how to personalise that skeleton into the usual suspects, like a general text editor, a simple drawing pad and a personal request, an OpenGL screen viewer.

John Gay, *via email*

Thanks for your suggestion – it's something that a few readers have requested before; and we've wanted to cover *KDevelop* for ages, but we just haven't been able to find enough



space! Come back next issue for the first instalment of our new *KDevelop* tutorial series.

There aren't any current books on KDE development out there, but as our star prize please accept *Applied C++* by Romanik/Muntz (Addison Wesley, ISBN 0-321-10894-9).

TSG's anti-competitive behaviour and suggest that others could do likewise or highlight the issue with their local office. The email address is neil.maxey@oft.gsi.gov.uk

My personal view is that with the Amiga we had Public Domain software, and with Linux we have Open Source; if commercial software companies had started threatening users of PD software, we would have laughed at them and told them to develop something better for their commercial dollars. If businesses had started using PD it would have been a marketing issue. I emailed links to the Halloween documents etc to DTI and OFT, so hopefully they can check out Microsoft actions in this issue also.

Mike Whelan, *via email*

There you go then. Will the OFT be Linux's saviour? In spite of the interest surrounding the case, it seems to have long since past being a vital concern to the average user.

Hit or Myth?

I was pleased to see the inclusion of *MythTV* on your recent disk. I have played with it before, but I didn't have the latest version – it sort of slipped my mind while I was playing with other things – but I was a little upset that their was no tutorial in the magazine. From my experience, running *MythTV* on a fairly standard PC (1.2Ghz P4, WinTV capture card) took a lot of work to get it working properly and even then, many parts – like the TV schedule – I couldn't get to work. It's nice that you include this stuff, but its not going to work for a lot of people unless you can explain how to use it...

Stephen Thompson, *via email*
Glad that you approve of LXF's choice of packages. As you can well image though, with the 100+ pieces of

Algol-lum

Contrary to Biagio Lucini's assertion, Fortran66 was NOT the first standardised programming language. This greybeard was programming in Algol58 and Algol60 (see www.masswerk.at/algol60/report.htm) back in the early 1960s.

Russ Herman, *Ontario, Canada*

I think we sort of covered this last issue, and of course you are right insofar as it goes. The reality is that FORTRAN was the first widely used (ie in America) standardised programming language.

The X86 factor

I've noticed that in recent months, the contents of the *Linux Format* coverdiscs have been almost exclusively x86-based. To make amends, how about a copy of Debian 3.0 m68k, my old Amiga is in desperate need of an upgrade!

Seriously though, whilst x86 is by far the most widely used of platforms, it seems sad to ignore the platform-independence of

GNU/Linux. How about a tutorial on installing Debian on an Amiga or an ST to redress the balance? Another good article that would be useful is how to set up remote dumb terminals via network or serial port and enabling dual-head graphics cards for additional workspace (they do say two heads are better than one)! Excuse the pun and keep up the good work!

Andrew Walker, *via email*

We have covered Mac hardware and even Amigas before (we had the install set for m68k debian on the CD many, many issues back), but with space limited, we do tend to concentrate on the most popular platform. We are intending to do a multiplatform DVD soon though, so keep scanning the disc contents over the coming months – there's no telling what goodies you might find!

SCO vs OFT

Regarding the *SCO v Linux* legal wrangle, I recently contacted the DTI, as I was concerned on how this could affect hundreds of British industries. We are uncompetitive enough already, without the threat of cheaper offshore file servers and business servers, for example, in Germany, TSG are not allowed to mention IP or issue threats against Linux users! After a few emails stating "HM Government has no

view on this issue," I

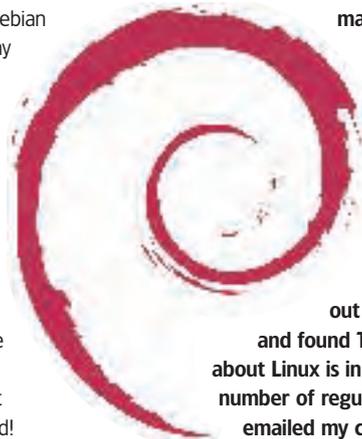
manage to obtain

an email

address of the officer in charge of this issue at the Office of Fair Trading.

I checked out the OFT site

and found TSG's comments about Linux is in breach of a number of regulations. I have emailed my concerns over



READER TIPS

CROSS OVER TO CROSSOVER!

I have dabbled and played with Linux, yet kept going back to that removable hard disk I had with Windows 2000 on it, because every so often I would come across a site that had something my browsers plug-ins couldn't cope with (*Mozilla* rocks, for those of you that ain't tried it yet).

Now I have seen the light – for a simple exchange of a few small pounds I now have all my favourite plug-ins using *Mozilla* on Linux: I cannot sing the praises of *CrossOver Plugins* enough, so instead I have decided I am going to not only use *CrossOver Plugins*, but also buy *CrossOver Office*. Good Work deserves praise and support, so I'm giving it in the small way that I can – by buying a copy of *CrossOver Office*.

Finally, for any Windows readers who are just testing the waters, I urge you to try *Firefox* or *Mozilla* instead of *Internet Explorer* – I bet you don't go back!
Eugene Meenan, *via email*

COPY UNPROTECTED

Imagine my dismay when I read in *LXF49* about this appalling new development multimedia wars, the use of CD copy protection. I'm not generally one for making copies of CDs, but when I'm travelling around I don't want to have to lug personal CD player plus box of CDs; not when I have a PDA and a CompactFlash card which is big enough to hold one or two MP3 albums (at 128K).

So the scene is set. After reading this article, I promptly received a CD for Christmas! What do I do? Throw it back and disown it? Hmm, not likely! With trepidation, I put the album on my old Kenwood stereo, which

predictably played it happily. So, girding my loins, I decided to try the same exercise on my Laptop.

The current spec is:

- Dell Inspiron 4000
- O/S Mandrake Linux 9.2 supplied by LXF
- Multimedia package XMMS (v 1.2.7)
- Kernel tweaked stock 2.4.22

Quite unexpectedly, the album played without a hitch! Finally – to prove a point and maybe a hint of nose-thumbing – I decided to rip it to 128K MP3 using *grip* (v3.0.7) utilising *Bladeenc* (latest version off the web). To my surprise, then I find that *grip* has done the business! The CD anti-copy protection method used was from "bmg-copycontrol". So I've either been very jammy, or this particular copy control method is dodgy!

All in all then, I've had a pretty good few days investigation, and I hope that this shows that all is not lost when the big bucks conglomerates get their

heads together! Incidentally, the note on the CD spine states that this method is NOT supported by anything other than Windows or regular stereo equipment. All I can say is "humph!"
Andrew, *via email*

THANKS FOR CSS

I've just received *LXF52*, and I saw in the *Mailserver* section a 'Thank You' from Derek Smith to one of your writers: Jono Bacon.

First I have to say a thanks myself, as I did find the series useful, and I use the final CSS file as a project 'prettifier' for things that I'm working on before I sit down and work out what colours I actually want – it makes working on things so much nicer! I'd also like to point yourselves and other readers to a small project that I found on Sourceforge, which is extremely useful and may have escaped your notice. It's called *CSSED* and the link is

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

This is a CSS editor in a similar vein to *TopStyle* on Windows. I have been using it for a while, and find it very useful. If the author is eventually able to add in all the features that he says he is going to, then I think it will be one of the most useful web tools available for Linux. Anyway, I hope you like it when you get around to checking it out.

Neil Williams, *via email*

DRIVERS

Great idea to put the drivers on the CD, however there will be lots of "tears before bedtime" with the Speedtouch drivers for anyone with the latest Speedtouch 330 Silver Revision 4 model. I have been struggling with this since getting it at Christmas: it will not run with the old drivers, and I was very frustrated with newsgroup posts – commonly, "I plugged it in and it ran!" Aargh! I have recently come across the new

modem-run and *mgmnt.o* etc files in a new Speedtouch rev 4 tar located at www.nemohackers.org/speedtouch.php.

Jim Riley, *via email*

KERNEL COMPILATION

In response to *Kernel Compilation*, *LXF51*, page 90, a forum poster expressed his desire to create his own customised distribution for academic use. I found the following info at www.distrowatch.com which your readers may find useful. Eagle Linux "is not a complete, ready-to-use Linux CD, but rather a document describing how to create one. As such, it is a great learning tool, suitable for educational establishments."

Additionally, while I don't want to discourage your endeavours, a lot of distributions seem to come and go for want of support. You could also lend your support to one of the existing distros targeted at academia. *Freeduc* and *Lorma Linux* are examples of these.

Chuck Coxhead, *via email*

MULTI-BOOT

RE: *Dual-boot duel*, *LXF52 Answers* page 94; why do you recommend booting into single-user mode for this problem? It seems a remarkably slow way of dealing with the problem. Why not leave it to the OS?

I have a multi-boot system with six operating systems. I have set the last parameter in each *fstab* line to 2 for all the partitions which the currently booted system uses – except for the root partition which should have 1. This causes the partitions to be checked during the boot process.

Then use *tune2fs -i* to set the number of reboots before checking is enabled. I use prime numbers between 19 and 47 to avoid having all partitions checked at the same time. For highly sensitive partitions, the value you use could always be much lower.

How you do this with *ReiserFS* is something that still eludes me...
Cecil Wallis, *via email*



MAILSERVER

« software that are included on the discs each and every issue, we'd be rather hard-pushed to include tutorials on all of them. The CD and DVDs files usually include whatever relevant documentation there is, though sadly, often there isn't much.

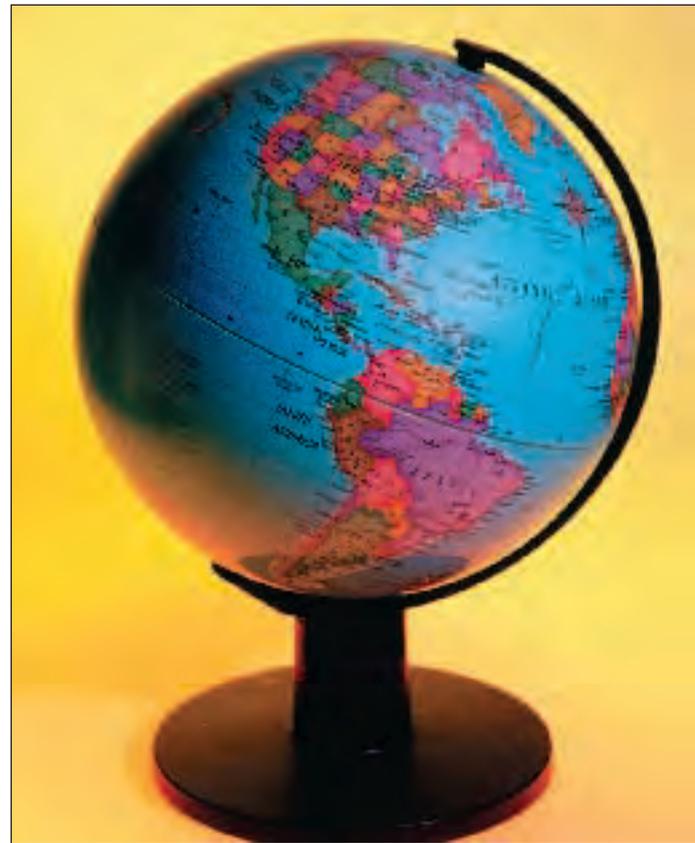
In the case of *MythTV* though, your prayers may be answered: we have a project idea for a *MythTV*-based PVR device, which should be appearing in the mag quite soon!

Geographical gaffe

In *LXF51* page 99, the geographic spot for Exeter LUG (73) should be moved approximately 1cm down the page to the other coast of South West England. At the moment the spot is approximately in the middle of Exmoor where I live!

Charles Woolmer, via email dictated using voice recognition software
Your house price is set to rocket - we have heard from a secret source inside the UK government that Exeter is being moved to your locality to secure it from the inevitable effects of sea levels rising due to global warming. This may seem like an extreme measure, but where else would the sons and daughters of the ruling elite go to University if they prove too dense for Oxbridge?

You'll be pleased to see that page 98 this month contains a fully



redesigned map, with all-new alphabetical LUG listings!

Why oh Why...

I dual-boot my PC between Linux and Windows, and I also like to keep

my hardware up-to-date: just splashed out on a new FX5950 card.

Though it's blisteringly fast under Windows, it was dog-slow in Linux; then I figured out that the drivers for Linux were whacked and

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

I needed to get some proper ones from nVIDIA. Now, this may just be another reason why Linux is really bad for games, but how come Windows was OK with the new card, even before I updated the drivers? And how come Linux (Fedora) is shipping without proper drivers? Until this gets sorted out I don't think many other people will treat Linux as more than a toy.

E.Rogers, via email

Well, there are a couple of things here. Firstly, nVIDIA does make its own drivers for Linux and Windows - one of the few manufacturers that does. Unfortunately, nVIDIA has never released the source to these drivers,

Helpdex

shane_collinge@yahoo.com



and probably won't, ever. This means that these drivers remain proprietary code, and as a result, they aren't likely to be shipped with many Linux distros.

Instead, most distros set up using the *nv* driver, which provides basic

support for these graphics cards, but currently this doesn't include any DRI or 3D support.

While obviously we would agree that having drivers included would be a good thing, the great majority of

current Linux users (ie those using it in business or IT departments) don't require 3D drivers for the latest gaming hardware, so while it may be annoying for home users, it isn't the end of the world for Linux in general. [LXF](#)

READER TIPS

MODEMS

For several months, I have been reading letters from readers anxious to know whether they will ever be able to get their Winmodems to work. I knew that I would eventually have this problem, as mine was just such. I bought one a year ago from Amazon. It was priced at £25.52 (net of carriage £5.01) and VAT but they do not quote that now. I checked a couple of days ago and they have one, similarly described at £29.99. The URL is http://www.amazon.co.uk/exec/obidos/tg/stores/detail-/electronics/B000089RS4/system-requirements/ref=ed_tec_dp_2_1/026-6329591-9500428

The box mine came in says that it is a Dynamode A220. A label on the internal card adds B-691 ver 1 02/05/31. The chips are Conexant. One says P4601-11 and E72060.1. The other Falcon 29 and 11627-31.

I found http://patrick.spacesurfer.com/linux_conexant_pci_adsl.html on the web. This is Patrick Mackinlay's site designed to make compatible Conexant PCI ADSL modems work under Linux. I had Mandrake 9.2 installed, but was not sure whether the correct options had been set when the kernel (supplied on the *LXF47* DVD (December 2003 issue) was built. Either way, it did not take too long to build a custom one and follow the instructions.

I received error messages during the build process. I emailed Patrick and he suggested that I use a 'C' compiler that had been tested with the modules downloaded from his site. The version of *gcc* on the *LXF* DVD as part of the package was 3.3; the latest one tested was 3.2.2. I downloaded



The Alcatel Speedtouch modem has received a lot of coverage in *LXF* of late...

and installed it and tried again – it all went smoothly and I carried on with the install. I executed the `test` command from a terminal. The modem seemed to be working.

Then, I tried to get *Konqueror* to find Google, but nothing seemed to happen. I should have rebooted because I later discovered that the relevant modules install during the boot process, and the modem comes up effortlessly. However, at the time I was uncertain of what was happening, so I dived into *LXF* to find that *Mozilla1.5* was flavour of the month. I downloaded it (under Windows) and rebooted under my custom kernel, installed and executed it. To my delight and amazement, up came *Mozilla's* home page – I was gobsmacked. I thought this might encourage others!

Alan Secker, *via email*

MAGIC

Thanks for reviewing *MagicPoint* again – it's long been a favourite of mine for running off quick presentations. A couple of handy hints you might also like to pass on to your readers are:

- The Tab feature, while often associated with bullet points, can

actually be used to set any text attributes. Thus, if you use two or three different styles of text often in a presentation, then you can assign them to tabs instead (the tab won't indent the text unless you specifically tell it to) and so you have a quick way for changing styles. (For instance, we project words for songs in our church, and often the chorus is centred italic text, with the rest in left-align normal, so I can set a single tab to mean centred-italic, and when I type the presentation, I can just tab in the lines of the chorus).

- You can set the *MagicPoint* window to stay windowed, and refresh itself from the presentation file, thus meaning that you can see 'in real-time' the results of changes you make to the presentation file. This is excellent when putting together a more complicated presentation – simply run the presentation in a smaller window (eg 640x480 on a 1024x768 display) and have it alongside your text editor of choice as you edit the file.

- *Vim* includes code-colouring for *MagicPoint* files, to make life a little easier, as does *Kate*. Ben Thorp, *via email*

PROTEUS COMPUTER SYSTEMS

www.proteus-linux.co.uk



We are the UK's No.1
for

Linux Computer Systems



3 year full warranty
FREE delivery!
Highest quality
GUARANTEED

Linux Compatible Hardware



BEYOND ADSL
PCI Modem
Only £14.99

Official Linux Distributions



SuSE Linux 9.0
Professional
Only £14.99

Linux Download Editions



Gentoo Linux
2004.0
3 CD Set
Only £6.00

Linux Books



"Linux in a Nutshell"
fourth edition
Only £25.65

More Computer Systems,
Software, Components &
Services available via
our Web Site!!

www.proteus-systems.co.uk

Tel: (01332) 518559

Fax: (01332) 510349



PROTEUS SYSTEMS LTD. 2004

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...

Mandrake 10.0

Six years is a long time in Linux, and this distro's come a very long way since it was just another Red Hat rehash **p18**

KDE 3.2

More evolution than revolution, but there are some great new apps folded into the mix **p20**

Unreal >> Tournament 2004

The top Linux gaming release of 2004 already? More fun than you can handle, if your GPU is up to the job! **p22**

PlanMaker 2004

Proprietary spreadsheet challenges *OpenOffice.org* **p26**

Subversion 1.0

Looking for a CVS replacement? Your search is over! **p27**

Book reviews

Top treeware for geeks **p29**



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.

MAJOR DISTRO

Mandrake Linux 10.0

Tom Wilkinson examines the ever-popular French distribution.



BUYER INFO

User-friendly distro aimed at the desktop market, and aimed at similar markets to Fedora Core and SUSE Linux.

- **DEVELOPER** Mandrake
- **PRICE** Free to download, boxed versions up to 199 Euros
- **WEB** <http://mandrakelinux.com/en/>

Mandrake started out life back in 1998 as what was – at the time – a slightly modified version of Red Hat. MandrakeSoft made a special effort to include a more user-friendly installer and a more ‘complete’ user experience than its main rivals, and its distro rapidly earned a reputation as a good start for new users.

Unlike competitors Red Hat and SUSE, MandrakeSoft allows download of its distribution in the form of ISO CD images as soon as it is released. This allows newer users an easy way to get started without much initial outlay – the SUSE installer requires a permanent network connection and Red Hat’s most recent distribution a fee of US\$179. This puts Mandrake in the position that it’s the one that many first-time users will come to, and it’s very suited to this position.

Installation

The Mandrake installer has for a long time been renowned for being easy-to-use and full featured.

One of the first things asked – before repartitioning or installation of packages – is what security level should be used on the machine. The lowest security allowed is ‘Basic’, for a standard desktop machine connected to the Internet, with each step clearly explained – increased security levels included checking system files and binaries for signs of tampering. The highest security level is for a firewall effectively blocking out the entirety of the Internet, including local networks, which you will then need to adjust to allow specific services.

Next came partitioning. Our test machine had come preinstalled with a

copy of Windows XP which filled all the available disk space. *DiskDrake*, Mandrake’s partitioning and mounting tool, made short work of this though, and offered the options of installing within the NTFS partition or resizing it. After choosing the resize option, the usual data loss warnings were thrown up, but the program coped admirably, even flagging that Windows should check the partition for validity the next time it booted. The default partition type is ext3. While *reiserfs* could also have been a good choice here, ext3 makes sense because there are no performance concerns and, should the need arise, it can be mounted as ext2 as part of a data recovery exercise.

Package selection included which desktops to install – KDE is already

selected and there are options for GNOME and *Window Maker* as well as the lesser-known desktops. There are also options for the types of application to install – console-based as well as graphical internet tools, and services such as *SMB* and the ever-popular *Apache* webserver.

One ‘feature’ introduced in 9.2 and still here is that advertisements are shown during the installation to hard disk (which took no more than about fifteen minutes in total), along with the usual messages thanking the user for installing. There were two that we noticed – one for Mandrake’s own MandrakeClub (which grants access to extras for the distribution – see page 109 for more) and one for a Russian anti-virus program. This is probably a

good way of subsidising free distributions, but in this case the advertised website was written only in Russian, so was perhaps not as widely appropriate!

Notable too, is that the paid-for distributions of Mandrake include the closed-source drivers necessary for using certain nVIDIA and ATI graphics cards, as well as the ubiquitous Speedtouch USB ADSL modem, easing the process of getting online.

The computer’s network connection was automatically set up at this point, obtaining the IP address, hostname and other information from a local DHCP server. There is also the option to set up a modem or ADSL connection using a local connection, available drivers allowing.

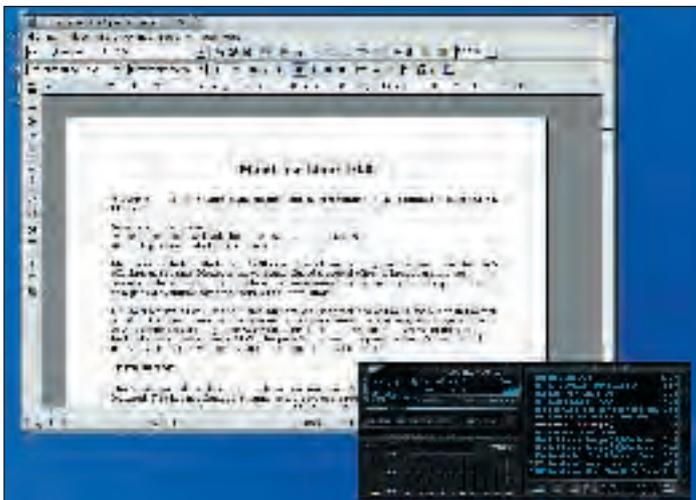


MDK Control Center makes it easy to make changes to the system, or install new software.

When it comes to the stage of adding users, it's possible to easily skip setting a root password. As in previous iterations of the distro, no warnings were given as to what effect this might have: this is rather disappointing, as setting a blank root password has some serious security implications that all users – new or otherwise – should be made aware of. Although a blank password does make day-to-day admin easier, it's totally insecure, as anyone with just a few weeks experience will know. One useful option here is the ability to automatically log in a non-root user at startup. On a single-user PC this can make for quite a time-saving option, though it's not advisable in multi-user environments.

If an Internet connection is available at this point, the machine will connect using it and look for updated packages. This ensures that all the software installed is patched against security problems before it even boots for the first time. Something like this is extremely necessary, particularly for machines which are to be connected to the Internet. It's surprising that not all operating system manufacturers – in the Linux sphere and elsewhere – make this option a standard.

On booting up for the first time, though, there were a few issues. The first and most obvious was that the monitor of our test machine hadn't been correctly detected: it was running at the correct resolution, but a very flickery refresh rate of 56Hz. In addition, the driver being used for the sound card was incorrect, resulting in a high-pitched screech coming from the speakers whenever a noise was supposed to be played.



OpenOffice.org is installed as the default office suite, offering solid interoperability with other productivity programs.

Each of these problems was easily resolved. In the case of the sound card, an alternative driver was available which worked correctly. The monitor was more difficult to fix though, and required the exact specifications to be found on the web. Apart from these niggles, though, the rest of the hardware worked fine, including the USB MP3 player which was automatically detected and its file system mounted.

Using Mandrake

On boot-up for the first time, the system offers users a choice between desktop environments: KDE, GNOME and Window Maker. Both KDE and GNOME have been customised by Mandrake to fit into their own desktop design, so there's really only personal preference to choose between them. We left this at the default – KDE.

The provided KDE desktop provides a familiar user experience to those users new to Linux – and on first boot included a screen offering links to the documentation as well as configuration options and a link to Mandrake's MandrakeExpert service. This provides an online free forum where users and Mandrake staff alike can offer solutions to problems you might experience. Unfortunately, clicking on the documentation option revealed that it wasn't, in fact, installed by default. There was an option to install it from the CDs provided before proceeding though.

A large number of applications are supplied with Mandrake – though not all are installed by default. Those which are include *OpenOffice.org* for word processing, spreadsheets, presentations and the like. We found



As the first distro to feature KDE 3.2 and a 2.6 series kernel as default, it underlines Mandrake's commitment to including the very latest software.

the lack of *KOffice* surprising; since it was designed to integrate better with the KDE desktop, it would probably be better suited as the default office suite. *OpenOffice.org* does, however, have the benefit of wider adoption, so it's easy to understand the decision to use it.

A wider variety of Internet applications were installed – there's a choice of three web browsers – *Mozilla*, *Konqueror*, and *Epiphany*, GNOME's default browser. Each of these works well, and in this case *Konqueror* is the default. Also supplied are a variety of email, news and FTP applications, with even more on the CD if the supplied ones aren't to your taste.

One thing that was surprising, though, was the lack of audio players installed. The default software for playing MP3s as installed was *Totem Movie Player*. This seemed to have some problems. Happily, *XMMMS* was included on the CDs, and this made a far better job of playing back audio. A DVD player was also not installed, but both *Xine* and *MPlayer* were supplied on the CDs – however, package selection is likely to be tightened up in the final release. Though the initially supplied application set was sparse, the distribution did make it easy to install additional software. The package management software – *RpmDrake* – lists all the packages currently in the database with options to search for specific names or keywords. The package database can be updated over the Internet, and if a copy is available on a local CD, it will prompt for that rather than download it over the Internet. This is a common-sense feature, particularly for dial-up modem users who would otherwise have to

wait a very long time (and possibly pay money) to install packages that have been downloaded from the Net.

The supplied Control Panel makes it easy to accomplish a number of complex tasks, including mounting remote SMB (Windows) and NFS (Unix) filesystems, and exporting directories so that other computers on the network – again running either OS – can see your files.

Including a 2.6 kernel rather than the more tried-and-trusted 2.4 version is a brave step for Mandrake to take, but a version of the older kernel is also available on the CDs in case of problems with certain hardware.

The distribution does, overall, succeed in its aims. It makes the installation of Linux exceptionally easy for the new user, while allowing the power and control that the more experienced user craves. There are still a few bugs to iron out, but given that our review copy was the last Release Candidate it's likely that obvious niggles will have been removed in the full product. Check it out for yourself – the coverdiscs this issue contain the full release of the Community Edition – see page 102 for essential install info. [LXF](#)

LINUX FORMAT VERDICT

FEATURES	9/10
INSTALLATION	9/10
EASE OF USE	8/10
DOCUMENTATION	8/10

A competently assembled distribution, new users could do a lot worse than to try Mandrake for their first foray into Linux.

RATING **9/10**



DESKTOP ENVIRONMENT

KDE 3.2

Jono Bacon reckons GNOME rivalry is good for both desktops.

BUYER INFO

Some prefer the Qt-based KDE, others the GTK-favouring GNOME, or there's always Window Maker...

- **DEVELOPER** The KDE team
- **PRICE** Free
- **WEB** <http://kde.org>

The seemingly tireless KDE team present a feature-laden release that is packed with new functionality. Not only have over 10,000 bugs been processed and the entire performance boosted, but KDE has been refined somewhat, and much of the clutter has been cut down. This has been applied to the entire desktop, menus, configuration and other places.

New features

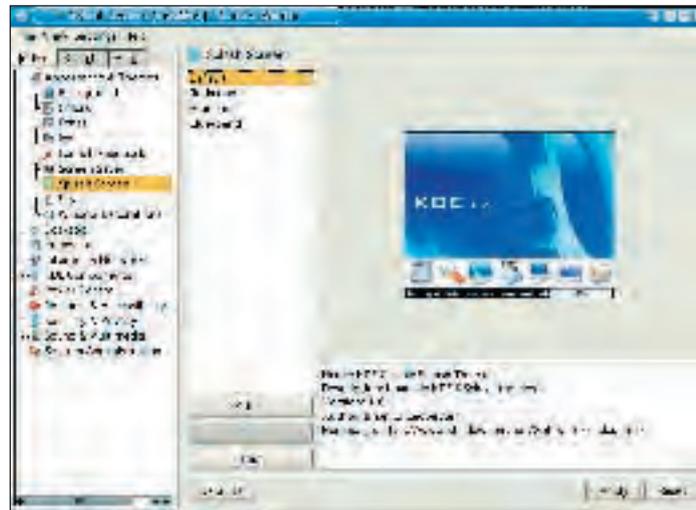
A number of new features have been added to KDE 3.2, many of which are subtle but useful improvements. An example of this is inline spell checking for web forms and refinements to *Konqueror's* sidebar. Speaking of *Konqueror*, the core KHTML component that is used to display web pages has had a number of contributions made to it by Apple. The reason for this is that KHTML is used in the Mac OS X Safari web browser. The updates and improvements have positively affected the quality of the HTML rendering and stability.

The KDE desktop itself has been improved with the new Plastik theme that is pleasant in its unobtrusiveness. A number of new icons have also been included with KDE at various sizes. The configurability for KDE remains ever present, and literally everything can be configured in terms of look-and-feel.

There are a few notable applications that are now part of the environment. The first is *KWallet*; a tool for storing passwords. This little tool is very handy, particularly in these days of the register-for-every-site Internet. Another notable addition is the *Kopete* instant messaging tool that has really set the standard for instant messengers. *Kontakt* is another tool



KDE can be visually spruced up to look any way you would like it to...



...as themes can be quickly and easily applied to your entire desktop.

worth a mention: it is the KDE equivalent of an *Outlook* style email/PIM tool. *Kontakt* can also work with the *Kroupware* groupware messaging server.

In Use

There is no doubt that KDE 3.2 feels more polished. Although most home users will probably not see a radical change in the desktop and its features, many will notice a performance improvement and the many subtle differences across the environment. Little changes such as the application switching menu and further compliance

with the free desktop project are littered across the desktop. Although you don't get any real initial oomph of surprise when you first install KDE 3.2, you will notice these little improvements in your day-to-day use of the system.

Throughout my use of KDE 3.2 I have found it to be incredibly stable. In some previous versions of KDE, I found some instability with certain parts of the desktop. Often, these problems seemed to be due to my configuration files getting their knickers in a right old twist. If I removed the configuration file, the problem often went away. These kinds of problems

don't seem to be as present in KDE 3.2. I encountered a few issues here and there, but these were often down to non-KDE-related issues.

Although KDE 3.2 is typically considered in terms of the desktop, there are of course a number of applications that are released with it. This includes *KDevelop*, *KOffice*, *Umbrello*, *Juk*, *KSVG*, *Kontakt* and others. Each of these tools and the many others have been updated to bring more features and increased performance to the system. *KDevelop* and *KOffice* are particularly notable. *Kontakt* is a 'meta app' that brings together *KMail*, *KOrganizer*, *KAddress* and *KNotes*: for an introduction to *KOrganizer*, see our last issue, *LXF52*.

Conclusion

The general feel of the environment is more polished, and the performance improvements are evident. It is great to see that all areas of KDE have been tightened and improved; we haven't yet noticed any areas that have been noticeably neglected.

The impression I got from this particular release is that the recent improvements in GNOME are really creating some friendly competition between the two environments. It is great to see this healthy rivalry, but it is also great to see how the two communities have come together with the free desktop project. I look forward to seeing how the *Project Utopia* (unified desktop window manager) and recent kernel improvements will improve removable device support in KDE and boost the continuing integration of the desktop. It is a bright future for all Linux desktop camps, and I am pleased to report that KDE 3.2 is another successful and impressive release. What are you waiting for? KDE 3.2 is on this month's coverdiscs – give it a try, and let us know what you think at the usual address! [LXF](#)

LINUX FORMAT VERDICT

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

Brings more integration, features, ease of use and compliance with open standards. There's *always* room for improvements in docs, not just in the Linux sphere though...

RATING **9/10**





LINUX GAME

Unreal Tournament 2004



The biggest game on Linux gets a refresh for the New Year: **Nick Veitch** and **Paul Hudson** get down to the serious business of blowing each other up...

BUYER INFO

GPU-intensive first-person shooter that steals the top Linux gaming crown from its predecessor.

- **PRODUCER** Epic Games
- **PRICE** £34.99

Spring means different things to different folks: naturists disrobe *en masse* with shouts of joy, wild animals wake from hibernation to procreate with abandon, and the LXF team gets early copies of the latest *Unreal Tournament* from Epic Games. Co-incidentally enough, we also happened to have our Ultimate Linux box sitting around; So, with headphones on, and fingers poised above keyboards, we got down to some, uh, *work*...

Install pain

We got the CD ROM version, and so were rather stunned to find that it comes on no fewer than 6 CDs, with a total hard disk space requirement of 5.5GBs. With each map weighing in at around 20MB each, it's clear that the geometry has really been ramped up to keep the pressure on modern GPUs. We had a few issues with Mandrake 10's supermount functionality playing

around during the install, so we ended up disabling supermount and doing the mounting/unmounting ourselves. If you have the choice, we strongly recommend you get the single DVD version and save yourself the hassle.

Shortcuts are automatically placed into your home directory for easy launching, and there's a Linux-specific readme to help get you started. One neat feature is that the installer automatically detects if you have 64-bit Linux installed and uses 64-bit binaries. Once installed, you do need your proper graphics card drivers installed, but that's par for the course really.

2004 vs 2003

If you played *UT2003*, you'll have an idea that Epic always uses the same engine for its games, but that it's under constant revision. Thus, *UT2004* shares much with its predecessor: rag doll physics, huge maps, and flak cannons are all in there, with tweaks.

There are now ten game types in total: Many have their own individual maps, but others (such as Deathmatch and Invasion) share maps to give you greater flexibility. Mutators can also be added to the mix to give an element of surprise and uniqueness to every game.

Assault mode has always been one of the unique aspects of this game, but Onslaught is now equally fun – it's essentially a modified Capture the Flag-type game where you win by chaining up mini-bases until you can finally take on the enemy HQ. The Bombing Run game is like basketball with guns: you need to carry a 'ball' through a hoop whilst avoiding fire from enemies; and there's also a new game called Invasion where you simply need to withstand assault from increasing hordes of enemies. As you can see, *UT2004* really does feel more like a sports event than a straight shooter – you can pick and choose your 'events' to fit the game style you enjoy the most.

Of course, the biggest improvement to the game is the addition of vehicles, and no, the developers haven't just added one or two little trucks: tanks, hovercraft, mechs, skimmers, HumVees, bombers, and many more all wait for you to climb in and wreak havoc. What's more, many of them – such as the tank – allow you rotate the turret independently of the main vehicle.

Graphically, this game is a dream to play. Thanks to the friendly chaps over at *PC Format* (who we later whipped resoundingly in Deathmatch –



REQUIREMENTS

MIN SYSTEM REQUIREMENTS:
1GHz CPU, GeForce 2 or higher, 5.5GB hard disk space, and 128MB RAM
RECOMMENDED SYSTEM REQUIREMENTS:
1.5GHz CPU, GeForce 3 with 64MB RAM, 5.5GB hard disk space, and 256MB RAM

hurrah for LXF), we were able to test out the same game running on Windows, and found it to be exactly the same – even down to the frame rates. This is very encouraging for Linux gaming, and shows yet again that Linux is more-than-ready to be a powerful gaming platform. There's free-form smoke, highly detailed models, an array of guns with fantastic light and sound effects, and explosions galore.

Of course, all these graphics are what has contributed to the high system specs: we'd recommend a 1GHz CPU and a GeForce 3 card or higher. If you're an ATI user, this you're looking at an 8500 or higher. So long as you have about 128MB of RAM, you should be OK, although having more does let the game pre-cache all the models for maximum performance – there's nothing worse than turning a corner and watching the game grind to a halt as it has to load your enemy's textures as he's shooting you.

EvilNick vs Hudzilla

Naturally the high-rolling game at LXF Towers was Nick (EvilNick) vs Paul (Hudzilla), which was largely induced by Nick issuing threats to give Paul a "crushing embarrassment" What we really liked was the large selection of one-on-one maps – super-small maps where there are only a few places to hide so that the action is always happening. And so EvilNick and Hudzilla met up, and it was quite a close match – see *Death Becomes You* for more.

Although the vehicles do add a new spin to the game, it'd be good if they were optional: Onslaught matches often become so focused on vehicles that it loses much of the whole first-

person fighting gameplay. Having said that, the vehicles aren't forced in the other modes, so again you can pick and choose what you'd rather play.

With so many games to choose from, you can spend a whole day just trying all the options out! After a while, we settled on the one-on-one maps, occasionally switching over to the Onslaught maps. However, if you don't have an Internet connection, relax – the bots are super-realistic, with slick animations, clear lines of thought, and even display cunning now and then. You can even jump into the same tank as a bot – while the bot drives, you can control the turret. Furthermore, there is a single-player tournament game where you have to work your way through qualifications and manage your team to success. In many ways this reminded us of the old Amiga classic by the Bitmap Brothers, *Speedball*, and we wouldn't be surprised to see a *Speedball*-esque event in *UT2005!*

If you want to see gaming on Linux succeed, it's time to put your money where your mouth is: snap this one up, and make sure you email Epic to let them know you bought *UT2004* just because it supports Linux... [LXF](#)

LINUX FORMAT VERDICT

FEATURES	10/10
PERFORMANCE	10/10
VALUE FOR MONEY	10/10

A top game all round: easily the most exciting entertainment release for Linux in the last 12 months – umpteen hours of fun for a very reasonable price!

RATING 10/10



Using the sniper scope to attack enemy positions in Onslaught mode is much safer than trying a half-baked WWI-style frontal assault.

DEATH BECOMES YOU

Settling LXF's professional grudges – with guns!

With Hudzilla being a 10-foot tall green monster and EvilNick playing as a skimpily dressed little girl (best not to ask), the two LXF gun-nuts set to playing head-to-head, mano a mano. After a few short minutes of gaming, it become

more apparent that the gaming session might be more useful as a display of the various ways EvilNick died, and so we proudly present to you some of the neat ways that you can vapourise your boss in *Unreal Tournament 2004...*



THE LEFT DIVE: EvilNick bravely attempts to dodge the barrage of rockets fired at him from six feet away, but ends up stricken with rigor mortis down his left side.



THE RIGHT DIVE: EvilNick has learnt his lesson about the left dive, and now dives to the right – sadly for him, with a similar end-result to the the left dive.



VAPED: A close range shot by Hudzilla's Flak Cannon left nothing but an EvilNick-shaped smear, which serves him right for trying a forward dive into the flak.



THE BLOOD BATH: Here, EvilNick claimed that Hudzilla was using "a special gun that killed him on the first shot." Again, it was the flak cannon that spelt his downfall.



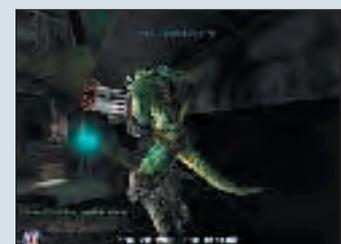
THE INVISIBUBBLE MAN: He may be invisible, but he's still as dead as kipper ties. From up high, the rocket launcher simply can't miss.



FLAMING DEATH: EvilNick must carry napalm in his pockets, but admittedly, a short-range rocket launcher blast is rarely survivable.



THE COMBO: EvilNick managed to display a majestic flaming death/bloodbath combo that made our sister mag *PC Answers* applaud.



CHAMPION TRIUMPHANT: Hudzilla does a victory lap with his trusty rocket launcher as EvilNick vows to stick with his day job.

ADVANCED SPREADSHEET

PlanMaker 2004 for Linux

Can a commercial Linux spreadsheet compete in the free market?

Roger Burton West thinks this one is in with a chance.

BUYER INFO

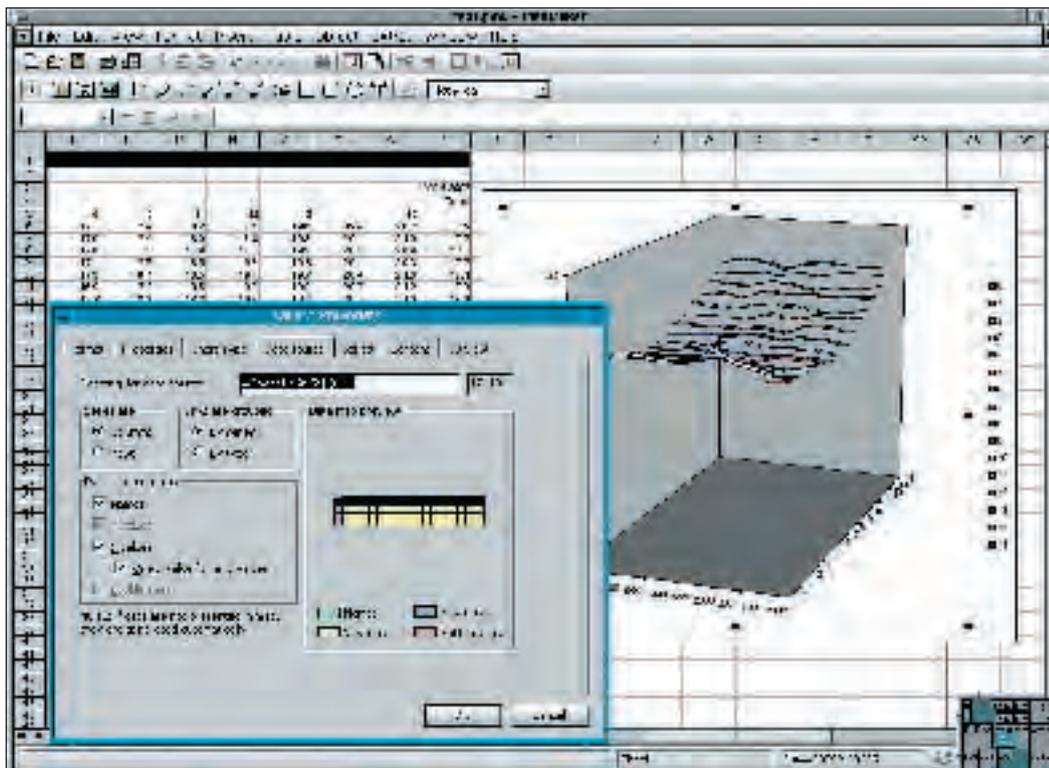
Another entrant into the Linux spreadsheet market already dominated by the likes of *Siag*, *OpenOffice.org*, *Oleo* and *Gnumeric*.

- **DEVELOPER** SoftMaker
- **PRICE** 49.95 Euros
- **WEB** www.softmaker.de/pm_en.htm

With its competitors mentioned above in the *Buyer Info* box all being free products, is there really scope for another Linux spreadsheet application, and conventional commercial software at that? *PlanMaker* is SoftMaker's companion product to its *TextMaker* word processor (released last year). SoftMaker clearly thinks it's enough of an improvement over the rest of the market to be worth the money; and, although it's not as far above the crowd as *TextMaker*, the company may well be right.

As with *TextMaker*, the program uses its own graphical interaction toolkit rather than relying on standard X widgets. For a cross-platform product, this is the cleanest coding solution, though it does mean the program stands out from other applications in the system. There are a few glitches – click-drag-release doesn't always work consistently for menus, for instance – but overall the interaction is handled well, with the usual suite of menus, toolbars, tooltips, and pop-up help. Moreover, most features are well-organised: graph types, for example, are all presented in a single list, rather than reflecting *OpenOffice.org*'s multi-layer structure. A reasonable selection of fonts is supplied, though the program does not use any others, which your X server might be prepared to offer.

Of course, as with *OpenOffice.org*, the program imitates *Microsoft Excel*. In the current corporate market, familiarity is probably worth more than features; but those who dislike the Excel-style interface will probably want to look elsewhere.



Setting up charts is fast and easy – like most tasks in *PlanMaker*.

Functionality

PlanMaker is certainly fast: on a low-end test system, it took around five seconds to start up the program and bring up a blank spreadsheet, as opposed to over fifty seconds for *OpenOffice.org 1.1*. Not entirely surprising, given that this is only a spreadsheet rather than a full office suite, but a complex spreadsheet was similarly faster to recalculate than under *OpenOffice.org*. The featureset seems reasonably complete; all the standard *Excel*-like functions are present and work as expected.

Layout features aren't exceptional, but their interface is well-handled; an array of well-chosen context menu items, as well as the app's responsive feel, makes setting up a fancy layout much faster than in other products. In fact, it is clear throughout the program that major effort has gone into making the user interface as efficient to use as possible. It may not suit every user, but it seems to work well for most.

The graphing interface is very clean, giving a three-step configuration (chart type, data range, and series identifiers); this is one area in which the designers clearly weren't slavishly following the standard model. In particular, when describing data ranges, the data layout is graphically represented as a backup to the usual text-based methods. Anyone who has spent hours tweaking a graph in *OpenOffice.org* will certainly find this vastly more convenient, and just as powerful, as that program.

Drawbacks

PlanMaker will read and write *Excel* files and the usual suite of little-used proprietary formats (no *Lotus* support, though), but it can't handle *OpenOffice.org* files. *SoftMaker* does itself no favours by fragmenting the market in this way: a company that's already using *OpenOffice.org* has no incentive to test *PlanMaker* in a mixed environment unless it has the budget

to obtain copies for everyone who might want to read a spreadsheet.

Some features are oddly placed or named (a symbol is something I "Insert", but a chart is an "Object"; and the 'Extras' menu seems to be a catch-all for the items that didn't fit elsewhere), and there are some stability concerns particularly in the charting modules, but the package is currently about as robust as *OpenOffice.org 1.0* and will doubtless continue to be improved. [LXF](#)

LINUX FORMAT VERDICT

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

A solid package, best at layout and graphics but doing a good, fast job across the board.

RATING 9/10



VERSION CONTROL SYSTEM

Subversion 1.0

**Freedom fighter Maurice Kelly takes control... by any means necessary!****BUYER INFO**

A replacement for CVS that runs on all modern flavours of Unix, Win32, BeOS, OS/2 and Mac OS X.

- **VERSION** 1.0.0
- **PRICE** Free
- **WEB** <http://subversion.tigris.org/>

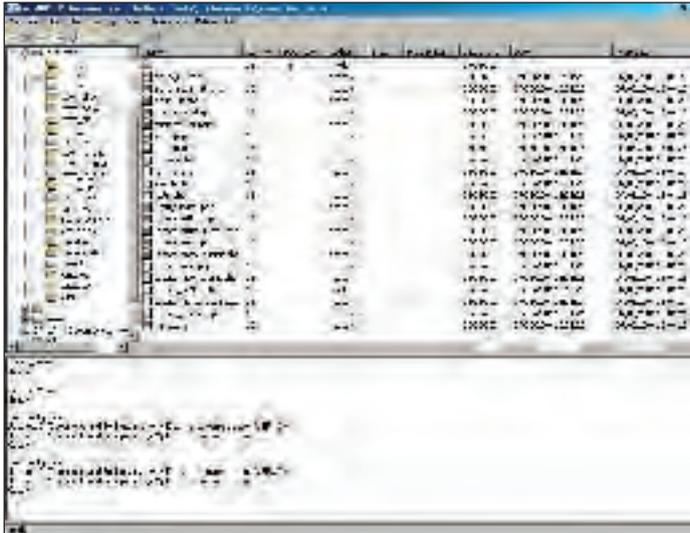
A version control system is a necessity for any large collaborative project.

Subversion is such a system and is the long-awaited successor to the widely used and popular *CVS* (*Concurrent Versions System*). Since Free/Open Source developers are – in many projects – geographically spread far and wide, *CVS* is a keystone tool for the FOSS community, yet even veteran users will acknowledge that it has some failings. Some of those failings were addressed by proprietary solutions such as Seapine Surround SCM, but now there is a stable Free alternative to *CVS*. *Subversion* was created back in 2000 and has recently hit the milestone 1.0.0 mark after a lot of hard work and waiting.

Contributed packages are available for some distributions and platforms on the project website; although we found none for our test distribution, we could opt to build from source. This was reasonably straightforward although the `INSTALL` file complicated things a bit more than necessary. After encountering problems hooking *Subversion* into an existing *Apache* installation we had to obtain and build one specifically for the purpose of running *Subversion*. Once up and running, it was quick and easy to manage. A sample project was imported and checked out by a user within seconds of starting *Apache*.

Apache?

Yes, you read correctly. A custom *WebDAV* module for *Apache* is supplied by *Subversion* and can be utilised to provide repository access over HTTP. Usefully, this provides access control via the established methods used by the HTTP protocol as well as additional security if HTTPS is enabled. It also makes a very handy



A number of third-party clients are available, the best probably being *RapidSVN* which is a cross-platform application built on *wxWindows*.

way to browse the repository directly through a web browser.

Local users can access a repository by way of the filesystem, and there is also a *svnserve* program which can be run as a daemon for those who want network access without the need to install and configure *Apache*. *svnserve* utilises a custom protocol which can also be tunnelled over *SSH* for additional security.

Subversion for CVS users

While it's not impossible that many users of *Subversion* will be users of other version control systems, it is fair to say that the majority will be existing *CVS* users. Sensibly, the *Subversion* team has elected not to make wholesale changes to the commands of the *cv*s client, so seasoned *CVS* users should be able to pick it up without a problem.

A major difference however, is the fact that each revision applies to the whole repository. Gone are the days when each file had its own revision level. Now we talk of "file as it appears in revision x" instead of "revision x of file" as the documentation is at pains to point out. As a user of both *CVS* and *Perforce* I found this strange to begin with, but by no means a major hurdle in getting to grips with *Subversion*.

Subversion is also somewhat safer than *CVS* – commits are atomic, so one file in conflict delays the entire commit. A conflict has to be specifically marked as being resolved before it can be re-submitted and conflict resolution is a breeze with the aid of a third-party diff/merge tool such as *tkdiff*.

Commendable effort has gone into making *Subversion* more resource-efficient. This can be easily identified when creating branches as *Subversion* creates a 'cheap' copy. Effectively, a branch is a copy of the original codeline and only changes are stored in the branch – this reduces the time taken to perform branches and offers a saving in disk space requirements.

Also new is the storage of a pristine copy of each source file inside the `.svn` directory (the *Subversion* equivalent to the *CVS* directory.) This allows a *Subversion* client to produce diffs at the client and send those across the network rather than complete files saving time and bandwidth resources when performing large commits. The local pristine copies also allow for some operations to be accomplished without network access. A binary difference algorithm makes the transference of binary files much more efficient than in *CVS*.

Migrating to *Subversion* would be a tough hurdle if there were no way to migrate your existing repository from *CVS*. Some tools are listed on the site, and *cv*s2*svn* is included in the tarball. Successful conversion isn't guaranteed though, and it's worth performing a test on a copy of your repository before adopting *Subversion*.

Third-party clients

The main *svn* client program is an implementation of the *Subversion* client APIs provided with the main package. The APIs are C libraries which can be incorporated into a third-party program such as an IDE or a dedicated *Subversion* client. Several third-party clients are already available including *RapidSVN*, and *ViewCVS* is reported to support *Subversion* repositories.

The highly modularised design of *Subversion* makes it possible for third-party developers to utilise any conceivable protocol to access a *Subversion* repository by implementing a repository access module on the server side and a corresponding client library. There is some very useful information about the architecture in the indispensable *Version Control With Subversion* – a free online book from <http://svnbook.red-bean.com/>

Conclusion

Subversion is one of those projects that has been well worth the '1.0 wait'. The feature-set is impressive and it is hard to find any real faults. I'm not going to be as happy working with *Perforce* as I used to be, and I certainly won't be going back to *CVS* unless I absolutely have to. **LXF**

LINUX FORMAT VERDICT

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

A worthy successor to *CVS* as the essential Open Source development tool.

RATING **9/10**



Red Hat Linux Bible Fedora and Enterprise Edition



Thou shalt worship no other Linux but Red Hat? Andy Hudson investigates...

BUYER INFO

- **AUTHOR** Christopher Negus
- **PUBLISHER** Wiley
- **ISBN** 0-764-54333-4
- **PRICE** £33.50
- **PRICE** 1065

Killing two birds with one stone, Wiley has released this *Bible* to cover both the Fedora Project and RH Enterprise Linux. To be fair, there isn't a lot of difference between the two flavours of Red Hat, and throughout this book there are clear explanations where the differences arise. My interest was in the Fedora Project and with my workstation ready and waiting I began to work through this book.

You can take two approaches with this book, both of which are entirely acceptable. You can read it from start to finish, working through each chapter and section or you can dip into it to find exactly what you need. The chapters have been well thought-out,



being fairly independent of each other yet the book as a whole really does flow very well. The other advantage is that it can be read by the newbie as well as the die-hard sysadmin, whilst not leaning either way. Whoever you are, you *will* learn new things from this book.

The usual stuff is covered – the history of Linux, installation, configuration and basic sysadmin. However, the depth that the book goes into is very impressive. For installation you can choose to install using the supplied CDs, FTP, NFS, HTTP or

Kickstart. For configuration, you can choose between GNOME or KDE. There's even a handy table on page 166 to the many Linux applications available for reference by those who are making the switch from a Windows environment.

It's fair to say that the book is split into two general sections designed for workstations and servers, divided by the chapter on understanding system administration. Sensibly this is the second largest chapter, with computer security taking a whopping and well deserved 70 pages to deal with.

Security is a topic that is fundamental to a good installation and it is pleasing to see the author give it such respect. Again, the choice is there between iptables and ipchains as well as information on intrusion detection, encryption and password protection.

The main theme that comes out of this book is that it gives you the choice to do things the way you feel comfortable. It neither warns against or recommends any particular tools but simply presents the options to you. For the seasoned sysadmin this is a great reference tool, but for the new Linux convert this is the missing manual for the Fedora Project and Red Hat Enterprise Linux.

LINUX FORMAT VERDICT

Brilliant book – far better than *Fedora For Dummies* from the same publisher – detailing pretty much everything you need to know about RHEL and Fedora.

RATING **10/10**



A Guide to Data Compression Methods

Paul Hudson reads a compressed sibling of the world's most famous compression reference...

BUYER INFO

- **AUTHOR** David Salomon
- **PUBLISHER** Springer
- **ISBN** 0-387-95260-8
- **PAGES** 295

As data compression experts go, the very mention of the name "David Salomon" is usually enough to hush any geek crowd. Not only is Salomon a renowned mathematician, physicist, and computer scientist, but he's also one of the leading lights in the compression community.

The primary reason for his fame is his book *Data Compression: The Complete Reference* (Springer Verlag; 3rd edition, ISBN 0-387-40697-2) which is due to be released any day now. First published in 1977, *The Complete Reference* is huge, and touches on every topic that's key to the compression industry. This *Guide*,



however, is clearly motivated by the need to produce a more lightweight tome: gone are the exercises, gone is much of the complex mathematics, and gone are the lengthy examples – this really does live up to the 'introduction' ideal.

Don't think for a moment it's easy, though. If you struggled at college-level mathematics this book probably isn't for you – within fifty pages you're already into stats tables and complex figures used to explain points made in the text. However, really this just shows how much of a good job Salomon has done at making the most of the limited space.

Although it might be a little tongue-in-cheek, I got the feeling he'd written a special lossless compression algorithm just for the book!

If you like details, don't worry – you'll find plenty discussed here, with the highlights being LZ77 and LZW, JPEG, and MPEG-1. As introductions go, this is more than enough to either pique your interest or turn you off for life. Understandably owing to its brevity, what you won't find is discussion of more advanced, modern formats such as MPEG-4 or JPEG 2000. With under 300 pages, there's only so much you

can cover. However, Salomon does his best to make up for any such detail shortcomings with a closing chapter on *Joining the Data Compression Community* – here there's a bundle to guide compression wishfuls in the right direction for more info.

If you've ever considered getting into the compression field, and you have a fairly good grasp of mathematics, this is a great place to start, and certainly worth buying before his larger and more expensive book – just in case you find it a little over your head!

LINUX FORMAT VERDICT

Won't help you if you're not mathematically minded or simply looking for ways to zip your email attachments, but excels as an intro to the nuts and bolts of compression.

RATING **8/10**



The Official Samba 3 How to and Reference guide

Nick Veitch reads the biggest Samba book ever.

BUYER INFO

- **PUBLISHER** Prentice Hall
- **AUTHORS** Terpstra, Vernooij and others
- **ISBN** 0-131-45355-6
- **PRICE** £27.99
- **PAGES** 704

This is actually a collaboration of many authors, which would be obvious to you if you could physically heft the weight of this book rather than just looking at the picture. The contributor list looks rather like a roll call of every notable *Samba* person you've ever heard of, with a few more thrown in for good measure.

Though collaborative efforts often tend to be disjointed, this work sets out its practical stall early on and continues in this same style throughout.

You might be amazed that there is so much to say about *Samba*: virtually



every facet of its operation, and interoperation with other systems, is explored, with numerous troubleshooting and fault-finding sections and snippets along the way. With practical worked examples of different configurations, it should certainly guide you to a working *Samba* setup, but perhaps more usefully will almost certainly have the answer to any problems you may

encounter in the future. There is some variation in style and tone in places. My particular favourite is a small section on using *MySQL* as an account backend, which begins with some statements suggesting that in the authors opinion this is a fairly bonkers idea, but then goes ahead and describes it anyway.

In all, 39 chapters cover every aspect of *Samba* you can possibly

think of, including Domains, Security, Printing, *Winbind*, character sets and using *SWAT*, the *Samba Web Administration Tool*. You can't even complain that the information is hard-to-find in such a large volume, as the book is laid out in a logical fashion – though the index merely borders on the adequate (seven pages doesn't seem nearly enough; when you try to search for something you'll see why).

This book obviously intends to be a complete practical reference, and while far from an entertaining read, it fulfils this purpose in a clear and unambiguous way.

LINUX FORMAT VERDICT

The best SAMBA book ever printed!

RATING 9/10



Learning Python, 2nd Edition

Snake-wrangling has never seemed easier for Nick Veitch.

BUYER INFO

- **PUBLISHER** O'Reilly
- **AUTHOR** mark Lutz, David Ascher
- **ISBN** 0-596-00281-5
- **PRICE** £24.95
- **PAGES** 596

The Python language predates Linux by a smidgeon. Originally written as a scripting language for the Amoeba distributed operating system, it was 'free software' right from the beginning. Python may have changed a lot over the years, but it is certainly more popular now, and is enjoying quite a renaissance due mainly to its cross platform characteristics and strong text handling abilities.

This book starts off well, with plenty of unintimidating entry-level stuff that sketches out the scope of the Python language, then actually manages to get



better, with plenty of detail and few unanswered questions. The scattered tips and warnings throughout are perhaps not as plentiful as they could be, but are nonetheless useful when they do put in an appearance.

One of the important things that often gets missed out of beginner's titles is to give some idea of not just how the syntax and language structure work, but also to give plenty of useful examples of what the language is

particularly good at. Fortunately this is addressed to a large extent in this book, with plenty of practical examples making important topics easy to understand.

This may be a beginner title, but toward the end of the book plenty of advanced topics are introduced (though not covered in detail), such as integration with *C/C++* code, *Jython* and various development environments.

The changes in this second edition are largely to do with updates to the

Python language itself. Although the current release wasn't available when much of the book was written, it is alluded to in enough detail where necessary. The updates for 2.2 that are included mainly address changes in classes and a few built-ins that were added to the language.

If you have read the previous edition of this book, you'll probably already know enough Python to migrate onto the more advanced *Programming Python* by Mark Lutz (O'Reilly, 2nd edition, ISBN 0-596-00085-5). If you are just starting out though, this is almost certainly the best title available.

LINUX FORMAT VERDICT

Simply the best beginners' title to the Python language so far.

RATING 9/10



Real World Software Configuration Management

Nick Veitch reckons there's more to SCM than just RCS and CVS.

BUYER INFO

- **PUBLISHER** Apress
- **AUTHOR** Sean Kenefick
- **ISBN** 1-590-59065-1
- **PRICE** £24.99
- **PAGES** 440

As the author suggests in the early pages of this book, Software Configuration Management is a role that can often be overlooked in the development and release of software projects. But it is a key role that brings together the work of all the other individuals and groups working on a project – it's not just about building installers. The first part of the book covers the role of SCM with an insight into the problems – mainly people-related – that come with the job.

Although there is a huge range of version control software available, only



two main contenders are covered in-depth – Microsoft's *SourceSafe* and the venerable *CVS*. A handful of other tools are discussed in the beginning of this section, and there is some helpful advice on choosing software from a commercial supplier, but the two systems outlined are the only ones covered in any great depth.

This is perhaps understandable, due to the depth that the book does go into, and the fact that these are

probably the most widely used software of this type. The *CVS* section covers anything you might reasonably want to know about deploying it for this purpose, and while you may find a separate work on *CVS* necessary, there is probably enough info here for most to get by on. Thought I'm not an expert on *SourceSafe*, this section seems to be equally as detailed.

By far the most interesting and useful section of the book for budding

developers is the final 200 or so pages, which discuss the marriage of the previous topics, based around the tasks you will need to complete. This is thoroughly practical and with plenty of useful advice.

Perhaps the biggest surprise is that what you might expect to be a rather dull and academic topic is covered in a refreshingly simple, conversational and often amusing way. The author has clearly suffered in an SCM role for some time, and many of the caveats and warnings that are passed on to the reader bear the undertones of bitter experience! **LXF**

LINUX FORMAT VERDICT

A thoroughly practical, and enjoyable book for both learning and reference.

RATING **8/10**



Sharp Zaurus SL-C750/C860 From Japan, Now Available In The UK



Amazing 640 x 480 Screen, Full Keyboard,
Word/Excel Compatible, Presentation,
PIM, Web Browser, Video, MP3 ...
Linux In Your ShirtPocket !

ShirtPocket
L I M I T E D

Tel: 0116 2701777

Email: salesdept@shirtpocket.co.uk
http://www.shirtpocket.co.uk



Roundup

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

C/C++ IDEs

Realising he is entirely incompatible with *Vi* and *Emacs*, **Jono Bacon** attempts to find solace with an Integrated Development Environment.

OUR SELECTION AT A GLANCE

- KDevelop
- Motor
- Anjuta
- MinGW Developer Studio
- WideStudio

The world of programming can be a frustrating one. Not only do you need to find the balance between donuts, coffee and actual programming, but you also need the right tools for the job, and I don't mean coffee-makers. The right tools can make a six-month job doable in four months or even less; the question is: which tools should you choose?

In this roundup, we have carefully selected five C and C++ IDEs and made them jump through hoops, walk on hot coals and otherwise test them to the fullest extent possible. Each is worthy of a try, but there are some clear winners that are definitely worth your time. Each has its own plus and negative points, and not all of these projects are heading in the same direction. They are all very individual

projects with individual goals and ambitions – it is therefore important to rate each of these tools on their individual merit.

It is likely that you know what you need to develop your code. You might use *Qt*, *GTK*, *wxWindows*, *ncurses* and write in C, C++, Java, BASIC, Python or something else. You will probably also need a debugger, source control system, profiler, documentation, patch viewer and even more donuts. The solution to this tool overload is by using an Integrated Development Environment (IDE) to tie everything together within a single interface. The key is to keep your mind on the code and not the tool.

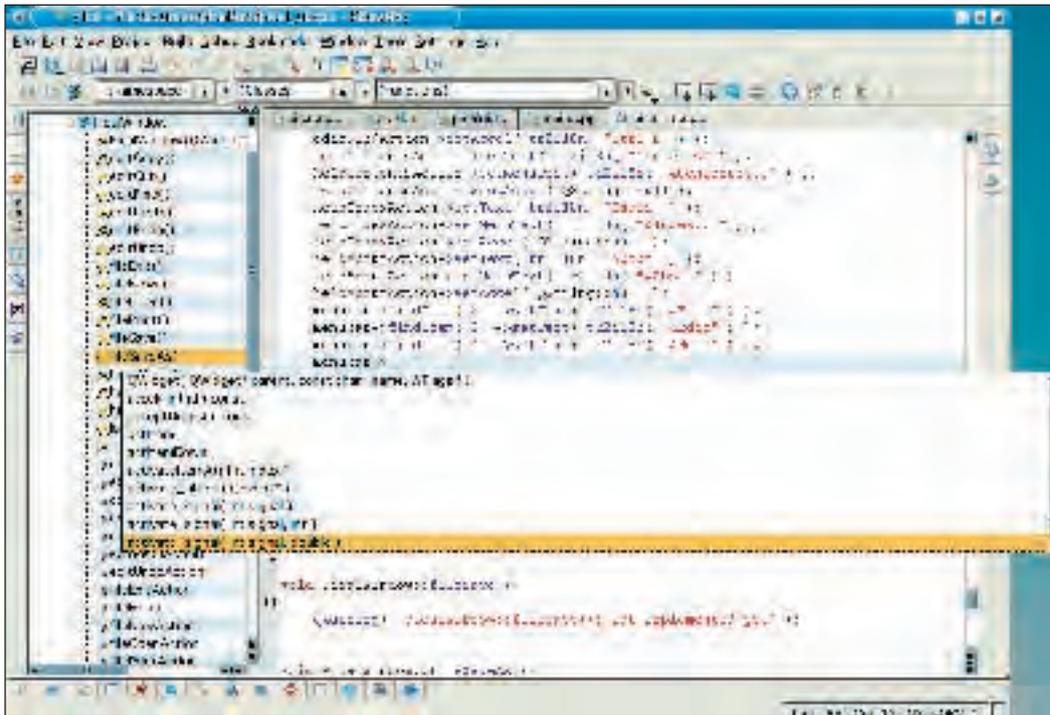
Unfortunately the delicate balance between a smooth, usable, productive IDE and a feature-rich, constantly

tweakable IDE is something that can quickly end up in hours of pressing Apply buttons instead of *creating* Apply buttons. When choosing the right IDE you need to make a choice up on what you practically need with your development and not necessarily what you think you will need.

Everyone has different requirements and views on what they expect from their software, and the

IDE space is one perplexed with debate and conjecture. You may well disagree with *LXF's* opinion on these apps: if so, we welcome the feedback at the addresses printed on page 12. The great thing about Free software though is having the chance to give them all a go. So, sit back, stretch your fingers, and prepare to get to know some tools that can turn your programming life from a chore to a joy.

'You need the right IDE for the job; the right tools can make a six-month job doable in four months or less; the question is: which IDE should you choose?'



A busy window packaged with functionality.

KDevelop

Code completion and KitchenSink all in one place.

■ VERSION 3.01 ■ WEB www.kdevelop.org/

KDevelop was one of the first projects to develop a fully integrated programming wonderland. Although an ambitious goal, the developers have worked hard to write, rewrite and refine *KDevelop* as a complete IDE. You don't just get C and C++ support here, you get pretty much everything else to boot.

The most recent versions of *KDevelop* (the 3.x series) are based on a completely new codebase that was rewritten to support extensibility and further additions to the application. The result of this effort has been a comprehensive jump in functionality and design; if you have used *KDevelop* 2.x or earlier, you will certainly notice this raft of improvements.

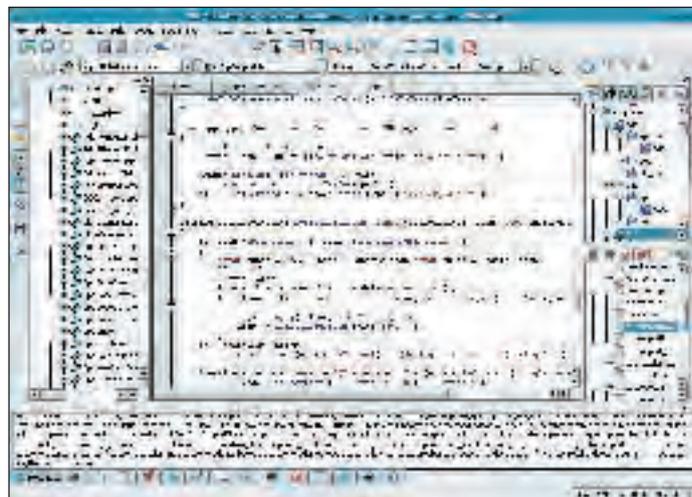
KDevelop packs a lot into its meaty KDE user interface. When you first load the application up you will be taken aback by the sheer amount of buttons, menus, widgets and other interface froth that is used in everyday development. A quick browse through the menus indicates that *KDevelop* can not only be used for developing code and compiling it, but source control, debugging, documentation, patching,

optimising and various other activities. In addition to these features, the multi-personality sidebar lets you constantly monitor your classes, files and other elements in your project.

Creating code

The start of a coding journey in *KDevelop* begins with its New Project wizard. This simple-to-use tool takes you

through the various steps of selecting a project, naming it, choosing if source control is used and finally generating a skeletal framework that can be moulded into your application. This project wizard is a great feature, not in the fact that it creates a framework for you (that has been the norm for years), but the sheer amount of different project types that are supported. For C and C++ programmers you can create KDE, GNOME, Qt, GTK and other types of project. This is a benefit in itself as it makes different platforms easier to get started with. Another useful aspect to your selected project type is that *KDevelop* adjusts its menus, sidebars and other features for that specific project. This means that



KDevelop 3.0 using KDE 3.2 and Plastik theme in IDEAL mode.

when you create a Qt *QMake*-based project, you will get the *QMake* manager to tweak the build system, but when you create KDE/GNOME project, you will get the appropriate Autom.

When you have created your project you are then ready to get started hacking your code. *KDevelop* supports a number of different editors that can be used within the environment and their associated keybindings (all of which can be configured). In addition to this, a perverse number of editor features are available within the settings dialogs. Literally everything can be configured from tab size, word wrapping, syntax highlighting, auto completion, line numbers and even different styles of auto-indenting. We are all aware of how heated the different *Vi* vs *Emacs* and other debates get in the Linux world, and it is no surprise that the *KDevelop* developers have chosen to try and accommodate everyone instead of locking certain groups of users out.

Highlights

KDevelop has a number of notable features that make it enjoyable to use. One of the main features is its support for code completion. This was one of the most requested features in the rewrite of the application and although *KDevelop* 2.x supported this feature, it only worked with methods in your own classes and not API library classes such as Qt. *KDevelop* now supports these API classes in a Persistent Class Store and it works remarkably well. As you might expect, this is all configurable. Another notable feature is the way *KDevelop* merges in a huge level of functionality by supporting external tools such as *Qt Designer*, *valgrind*, *doxygen*, CVS and others. The developers have done a great job here and not only prevented reinventing the wheel, but actually made the integration feel...integrated.

LINUX FORMAT VERDICT

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

KDevelop itself is GPL, but for commercial development using some libraries, you will need a professional licence; for instance, KDE/Qt apps need a licence from Trolltech.

RATING 9/10



Motor

Not everyone wants to get all GUI.

■ **VERSION** 3.2.4 ■ **WEB** <http://konst.org.ua/motor/>

Within the C and C++ programming community, there are a great many people who have never even heard of *Motor*. This is no doubt due to the largely GUI dominated world of IDEs; a world that *Motor* is not a part of. Yes folks, you guessed it: *Motor* is a console C/C++ IDE.

It is quite clear from the offset that *Motor* is targeted to console mode applications as opposed to GUI applications. This is in many ways its greatest strength. If you are developing a fully console mode tool, there often seems little point in using the performance lag of a graphical interface to develop your code. There is a good possibility that you could combine *Motor*, *Links* and other terminals to produce your tool and cut down on this bloat. *Motor* also has a slight bonus in the fact that it is tiny. The sheer compactness of *Motor* aids faster performance and can open up C/C++ development on older boxes previously only considered as bin-fodder.

Features

Considering *Motor*'s compact size, it packs a surprising amount of

functionality under the bonnet. First and foremost, *Motor* gives you the ability to generate a C or C++ program with a customised build system, or by using the popular *automake/autoconf* system that is utilised by most free software projects. While configuring your project details you can also add the standard GNU documentation files (INSTALL, README etc.) and make use of *gettext* to enable internationalisation within your project. When these settings are complete, you can either have your source code generated, or start from scratch.

When *Motor* is not in project creation mode, it falls into its most familiar of guises – editing mode. When I say modes in this context by the way, don't assume that I am referring to a *Vi*-esque twin-mode way of working. *Motor* puts you into the main editor view and you can begin hacking your code with full syntax highlighting, code searching, search/replace, block indent and other features. When you need to switch to another file, you just load up the file list and select it. *Motor* also allows you to compile your project within the

environment itself instead of shelling out to a different console to run make.

There are two features that I was surprised to see in *Motor* when I started using it. The first is full version control management with CVS. *Motor* has managed to squeeze in the features in CVS that you need to use 90% of the time. This includes importing code, making updates, checking out, tagging releases and more. Naturally, the level of support for version control is not what you would expect in a full CVS client, but nonetheless, it is well integrated. The second feature that I was very impressed to see was an integrated debugger. This built-in bug hunter is actually quite impressive, and there is a good level of functionality included. Notable debugging features include loading core dumps, adding watches, setting breakpoints, line-by-line debugging with steps, and more.

In use

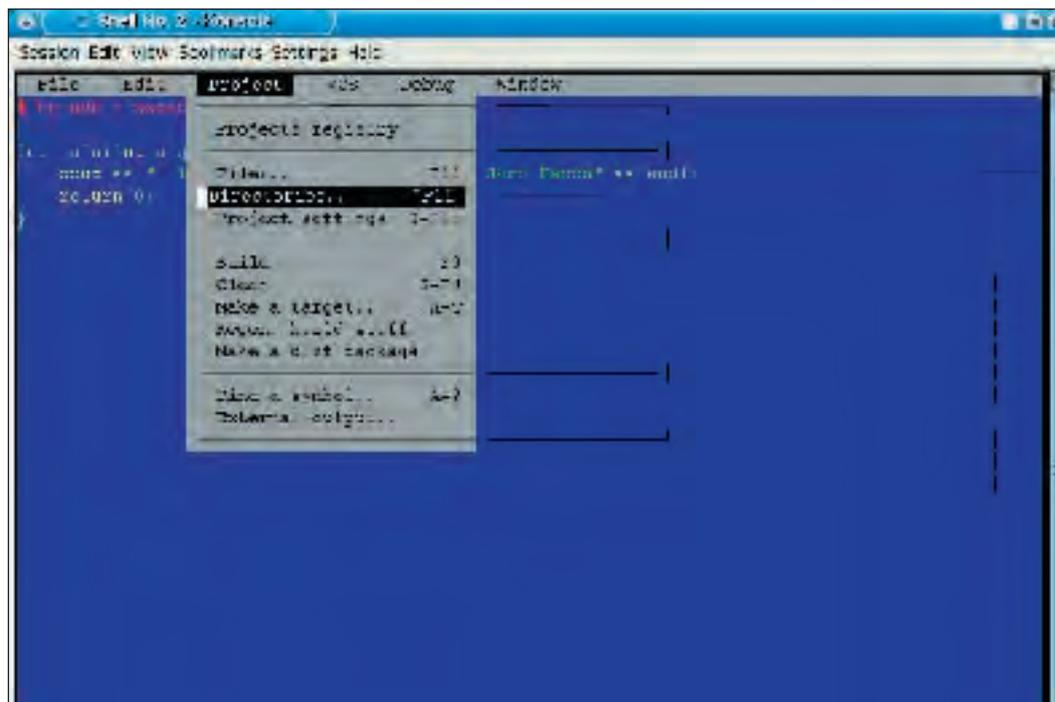
Motor is a great little IDE; there is no doubt that *Motor* does not even scratch the surface in terms of the functionality offered by tools such as *KDevelop* and *Anjuta*, but it doesn't set out to. In general usage, *Motor* seems fairly content being a fairly expansive editor with some IDE-type functionality thrown in to make it better for writing C/C++ applications.

When I was using it, one area where *Motor* really took me aback was

in its ease-of-use and polish. From my experience as a Linux user, I must have been through every editor and every IDE there is on offer. Out of the mountain of console editors, I find that *Motor* actually holds its own. This problem of finding a simple-to-use yet flexible editor has certainly not plagued just myself; you don't have to go very far with Google to find other examples! *Motor* successfully presents a simple-to-use editor that will let you perform most editing functions fairly easily. This is certainly no *Emacs* or *Vi* though: if you are expecting everything but the kitchen sink, you will be sorely disappointed, but if you want a simple editor that works, I think that most users will be impressed.

While using *Motor*, there were only a few things that I really wanted that were left out. The first is without a doubt a class parser. When writing in C++, a class parser is like a roadmap to your tangled web of code; without it, I must admit I did feel a little lost. This leads me onto to my second request: I would also like to see a windowing structure within *Motor* so I could, for instance, have open a class parser, two files and the messages window. The fact that I can only edit one file at a time is a big let-down with the IDE. With any C and C++ coding, it seems fairly common that the coder has a particular .h and .c/.cpp file open most of the time. If I am editing *MyClass* for example, it is generally useful to have both the *MyClass.h* and .cpp files open. One final request that would be nice to see in *Motor* is code completion. It would be great to be able to easily access methods from relevant objects in my code.

One other point that should be noted about *Motor* is its quality as merely an editor. If you are also looking for a simple-to-use editor and just cannot get on with *Vi* and *Emacs*, give *Motor* a whirl.



Content to be nestled in a terminal, *Motor* suffers no GUI glitter.

LINUX FORMAT VERDICT

FEATURES	6/10
EASE OF USE	9/10
DOCUMENTATION	5/10
PERFORMANCE	10/10

Not an IDE for the inexperienced, as there's no comfortable GUI to click around. If you need a lightweight IDE/text editor though, it's quick and simple.

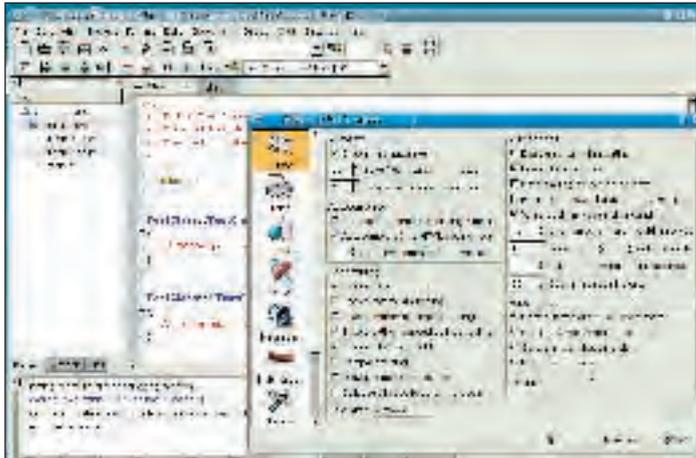
RATING **8/10**



Anjuta

There is no place like GNOME...

■ **VERSION** 1.2.0 ■ **WEB** <http://anjuta.sourceforge.org/>



A lean and easy to use interface is presented with Anjuta.

Anjuta is a GNOME/GTK based IDE that supports development with C, C++ and other languages, a very strong IDE in terms of integration and use. The environment looks similar to most other IDEs – sidebar, main editing area and compilation messages pane in a clean

and fluid interface with a great selection of tools and facilities. One thing to be aware of though is that this IDE is most primarily aimed at GNOME/GTK development – there is no KDE support but there is support for console, terminal, *wxWindows*, *xlib* and *X Dock*

applications however. Within the sphere of GNOME/GTK development however, *Anjuta* is certainly a promising tool with GTK, GNOME, *libglade*, *gtkmm* and *bonobo* targets among others.

In use

When you use the project wizard, the sidebar will fill with content and you can begin hacking your code. This configuration typically begins with the editor, supporting syntax highlighting, code folding, automatic formatting, line numbers and much more. The editor can also have its keybindings reconfigured, and even strange options such as the width of the editor caret (the cursor) is configurable.

The actual build process of *Anjuta* generally worked quite well, although I did experience a few flaws in the system when initially generating a project and making it configure. However, a bit of knowledge of the overall Linux build process should help you to resolve small problems such as this. The only real grudge is that the error messages in *Anjuta* were not particularly helpful when the problem occurred. This could be largely the fault of the build tools that *Anjuta* relies on, so the finger of blame is not entirely pointed at *Anjuta*, but it'd be good if this issue was ironed out.

Anjuta comes into its own when it is combined with the excellent *Glade* GUI development tool. *Glade* is the equivalent to what *Qt Designer* does for KDE/Qt development, producing a GTK interface in a WYSIWYG fashion. *Glade* can be combined with *Anjuta* by running the two apps separately, and this allows for a more RAD-orientated development process. Although the two apps work well individually, there is little integration within *Anjuta* itself for *Glade* forms, but detracts only slightly from the process: not much of an issue, as you simply save your *Glade* forms and then load them up in *Anjuta* afterwards. Better integration appears to be planned for the future, and I look forward to this update.

LINUX FORMAT VERDICT

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

To many, *Anjuta* is the equivalent app for the GNOME desktop that KDevelop is for KDE, but doesn't have as much out-of-the-box functionality as some of its rivals.

RATING **8/10**





World class support that never closes

Every Fasthosts customer enjoys full access to Xtreme Support: unsurpassed customer service, faster turnaround for technical problems and an unparalleled level of help and advice.

- 24 x 7, 365 days telephone, email & online support
- Xtreme Support website - with knowledge base, FAQs and more
- Trouble ticket system for a fast accurate technical response

With multiple internet connections and our advanced UK data centre, it's no wonder we're known as **'the power behind web hosting'**.

Linux Apache PHP hosting

Home Pro account

Email virus scanning, web mail, SMS email alerts and 1 GB of web space on reliable, Linux servers, make this the ideal account for great value and great features. With support for Perl, PHP and CGI and optional MySQL databases, Home Pro gives you everything you need to create high spec, feature rich websites.

Also includes:

- Matrix Stats - real time data on visits to your site
- 300 POP3 email and forwarding addresses
- Easy to use website builder
- Access to URL security, streaming media and much more
- Manage everything from the web with our award winning control panel

FREE INSTANT SET-UP

£8.99

Monthly +VAT or £70 annually

All your hosting needs

- Business packages from £14.99
- Developer packages from £19.99
- Unlimited Reseller package from £50
- Dedicated server solutions from £89
- UK's best value domain names at www.ukreg.com




For instant advice and set up, call

0870 888 3833

or for more details and to launch your service within minutes, visit www.fasthosts.co.uk

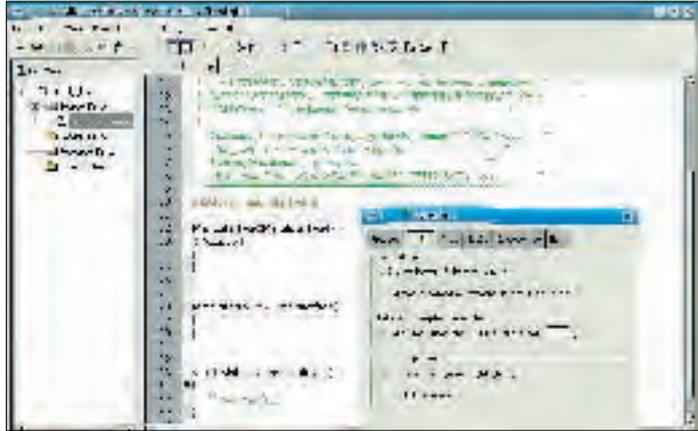


www.fasthosts.co.uk

MinGW Developer Studio

Is MinGW a minger or a winner?

■ **VERSION** 0.98 ■ **WEB** www.parinya.ca/home.html



MinGW presents a simple interface and sticks to it.

On first sight, MinGW appears to be fairly comprehensive, built using the wxWindows toolkit and on its website claims the following: "MinGW Developer Studio is a Cross-Platform C/C++ IDE

(Integrated Development Environment for GNU GCC Compiler system". Most other IDEs are only usable for Linux and Linux-style systems. wxWindows is fairly well established as a cross-

platform toolkit and I was interested in seeing how MinGW performed.

MinGW comes pre-packaged with the fairly typical range of IDE fodder that you would expect for a minimum level of functionality. This includes support for different projects such as GTK/terminal/wxWindows/GNOME/library targets as well as build support, multiple views etc. The editor built into MinGW supports syntax highlighting, auto indent, line numbers and smart indent. Sadly, in terms of editor features, that's about it: no font configuration, code folding, key binding configuration or other common advanced tweaking. As with other Free software projects, this is probably being worked on, so it is worth keeping an eye on the project.

While using MinGW I did spot two noticeable features: the debugger and the active configuration. The debugger worked as expected and is very similar to the old debugger in KDevelop 2.x. The active configuration tool lets you

create different configurations for different development periods such as debug and release.

All in all, very little of MinGW actually worked for me. The interface seemed to be rather unintuitive at times, and the IDE did not generate code when I created a new project: I assume the code needs to be manually created. I also had difficulty working with a project and the editing and compiling process seemed flaky at times. When using the IDE, I really got the feeling that this will be a great IDE in six months' time when some more dedicated development has occurred.

LINUX FORMAT VERDICT

FEATURES	5/10
EASE OF USE	8/10
DOCUMENTATION	5/10
PERFORMANCE	8/10

We spent more time messing around with the IDE than hacking code. Glade support and a more intuitive GUI would greatly help matters. Has some potential though...

RATING **6/10**



WideStudio

RAD: just a fad, or really bad?

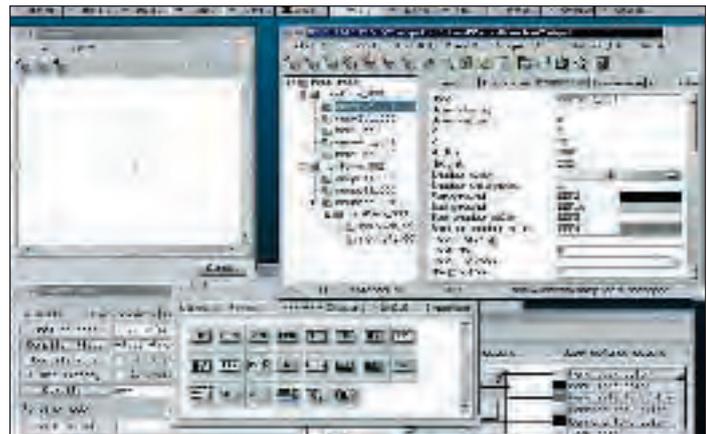
■ **VERSION** 3.50-2 ■ **WEB:** www.widestudio.org/EE/index.html

This is a different kind of IDE in the sense that it is specifically geared towards writing Rapid Application Development (RAD) C/C++ and Java applications, largely due to its built-in form designer, procedure editor and other features. WideStudio is not your typical IDE that will do the hard work of generating a project; its aim is to generate a project so that you can visually develop your app and attach relevant code to its functionality. In RAD, this is common ground, but typically higher-level languages (such as Visual Basic or REALBasic) have been used for the code. Using C/C++ and Java brings the concept to a different level.

The idea behind WideStudio is not too dissimilar to Qt Designer/ Glade but seems to be more integrated. With Qt Designer/Glade, you design your interface in a separate app, generate the code and then hack your real code. With WideStudio, it all seems to

be thrown in together, and you can do most of these things within the app.

WideStudio includes a main app builder tool and a suite of utilities to tune projects. The main app builder lets you create a new project, add windows and design your app layout. When creating interfaces, you have a good selection of widgets such as buttons, textboxes, scroll bars etc. In addition to common widgets, facilities such as specialised drawing widgets, different window types, non-GUI facilities (such as database, networking, source control etc) can be edited in a not dissimilar way to the Qt Designer/ Glade equivalents. As such, WideStudio essentially ties you right into its widget set, so everyone coming to WideStudio will need to go through a learning curve to learn how to use its custom widget set. It is a real shame that the app doesn't support Qt/GTK/wxWindows as this could have been a really interesting and widespread tool.



A rather dated look by today's standards, but WideStudio packs a punch.

I was impressed with how simple it is to create a GUI app. I had a simple window designed and implemented within 10 minutes with no coding. This in itself would be a useful prototyping tool. WideStudio is fairly intuitive, and making links between widgets and functionality was very simple; much of the code was generated for you.

WideStudio has been carefully crafted in turns of creating interfaces and connections, but the true test is in writing the code. To fully appreciate and test WideStudio I would need to learn an entirely new widget set and I simply did

not have the time needed to cover this information in sufficient depth.

LINUX FORMAT VERDICT

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

If you can live with its custom widget set, this RAD/IDE could be handy. Giving it a try is the only way to find out...

RATING **7/10**



C++ IDEs THE VERDICT

All of our five contenders have put up a worth fight to be crowned the king of IDEs in this roundup, but there is a clear winner: *KDevelop*.

KDevelop is a great IDE; many readers will be aware that I have been involved with the KDE project, and you may feel that I could be biased, but the true merit of *KDevelop* is present in itself. The *KDevelop* team has done a superb job of creating an IDE that is extremely flexible, unobtrusive and a joy to use. I cannot think of anything else that I would possibly need in an IDE. If there is any advice I would offer the developers though, it would be to look into the best methods of structuring the interface. *KDevelop* is a complex and feature-laden application, and as such, it needs to be carefully crafted so as not to be confusing to the new user. Overall though, a great tool for C and C++ programming.

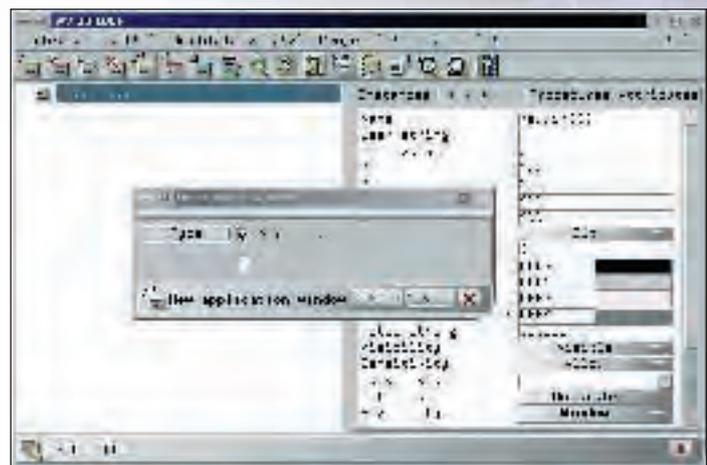
As an IDE, *Anjuta* works well but when I was using it I couldn't help but feel the true potential of this tool has not been realised yet. The reason for this slight cynicism is that the actual interface and programming tools in *Anjuta* are really good, but it is a shame that the application is limited to building GNOME/GTK/*wxWindows*/terminal applications only. If the developers were to open their scope for the IDE a little more, things could really get exciting. I could imagine *Anjuta* taking on a similar role that

KDevelop has mustered and support a huge range of different languages and platforms. Although this could well happen in the future, I do have a contradictory respect for the *Anjuta* developers for keeping the focus on the included platforms. Definitely one to keep an eye on.

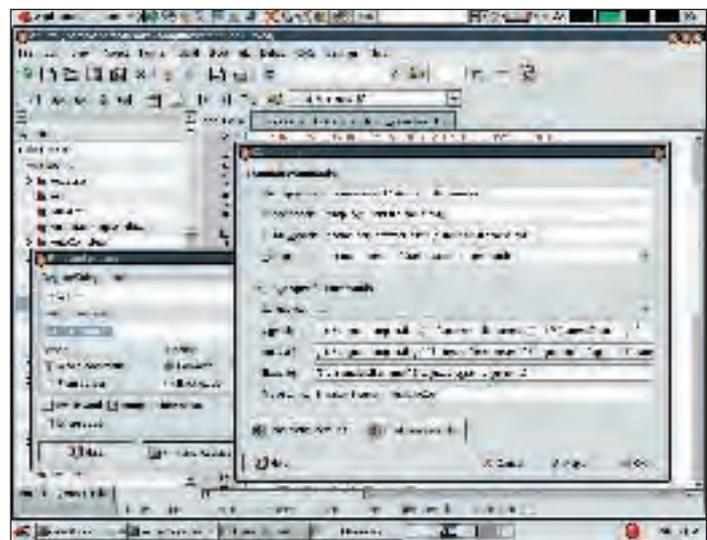
If you are someone who needs to write console applications, *Motor* is a great choice and at least worthy of a test. Despite its more economical stance in terms of features, this economy is passed on in performance, a factor that makes programming on older machines far more feasible. *Motor* provides what I would consider the basic level of functionality for an IDE, with a few niceties rolled in such as CVS and debugging support. This range of features has all been encompassed into an easy-to-use interface and the whole shebang worked flawlessly for me.

WideStudio is a very interesting and functional RAD IDE that is sadly let down by its lack of support for Qt, GTK, *wxWindows* and other widget sets. Five years ago, *WideStudio* could have made a real impact, but the *Qt Designer* and *Glade* developers have stepped up to support this RAD functionality in existing IDEs and widget sets.

MinGW Developer Studio is a simple IDE in terms of features but not really usage. With a little user interface tweaking in the right direction and some features such as code generation,



WideStudio: a great RAD app, but still playing catch-up in IDE terms.



Customisable commands could make Anjuta a future Linux killer app?

TABLE OF FEATURES: PLATFORM SUPPORT

IDE	C Support	C++ Support	Qt	KDE	GTK	GNOME	Bonobo	Terminal
<i>KDevelop</i>	yes	yes	yes	yes	yes	yes	yes	yes
<i>Anjuta</i>	yes	yes	no	no	yes	yes	yes	yes
<i>Motor</i>	yes	yes	no	no	no	no	no	yes
<i>WideStudio</i>	yes	yes	custom	custom	custom	custom	custom	custom
<i>MinGW DS</i>	yes	yes	no	no	yes	no	no	yes

code folding, CVS/Subversion support and font adjustment, *MinGW Developer Studio* could really bolster its own potential. *MinGW Developer Studio* is currently free to download but not Open Source; I really urge its author Parinya Thipchart to Open Source this IDE and further pave the way for its true potential. [LXF](#)

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 *Mailservers* pages.

HOT PICKS AT A GLANCE

GCompris	36
SegusoLand	37
KolourPaint	37
HotSaNIC	38
FileGarden	38
Bookcase	39
Gnofract 4D	39
Sanitizer	40
GCrontab	40
Senken	41
Blobwars	41

HOT PICKS AWARD

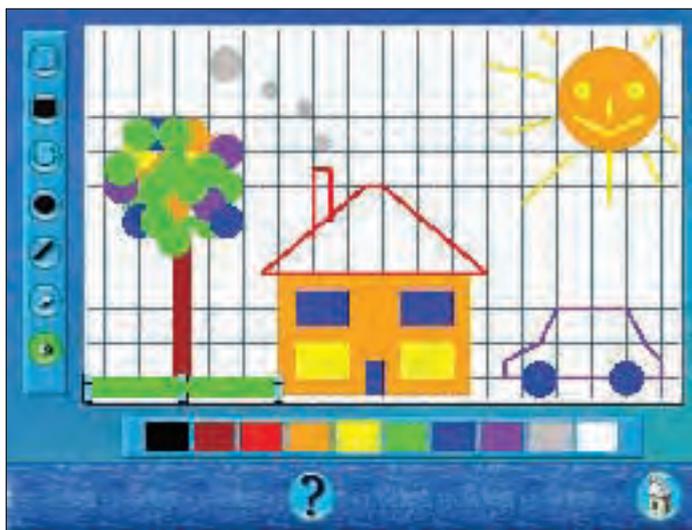
Everything covered in our *Hot Picks* section is unmissable, but every month we'll be singling out one project for outstanding brilliance. Only the very best will be chosen!



EDUCATIONAL SOFTWARE

GCompris

■ VERSION 5.2 ■ WEB www.ofset.org/gcompris/



Diane? The unhappy flower doesn't think much of our image recognition...

Edutainment, like Multimedia, is one of those annoyingly vague terms dreamt up by buzzword-happy marketing managers in a fit of rampant boss-impressing 'proactive initiative'. Hurling onto the masses by Mindscape's *Mario is Missing* and the equally enjoyable *Where in Time is Carmen Sandiego*, edutainment on the 16-bit consoles worked fairly well, combining traditional games with educational elements. There are a number of similar projects for Linux and since it was featured on last issue's discs, we've had plenty of requests from readers to give the superb *GCompris* more coverage.

GCompris (from 'j'ai compris', or 'I Have Understood') is a compilation of games and tasks designed to entertain a child while he or she learns. Although the target age range will vary depending on the task in hand and developmental stage of the child, *GCompris's* authors suggest 3 years and upwards is best.

Usefully, *GCompris* can be installed into a single directory, making for easy

removal and addition of other components later on. It also makes it quick to back up, so if it crashes or a child corrupts it, you can simply copy it back over. Building from source is a doddle as it follows the usual process; you'll need to have the libbasetml library installed first, and once compiled you may have to specify the **-A <sound>** flag when launching (eg **-A oss**).

The *GCompris* interface is colourful and laden with cute touches: the default blue skin can be switched to pink for girl children, the icons are smooth and mostly self-explanatory, a chunky red mouse pointer and smiley/sad faces for feedback, the mini Help guide finishes off an impressive design that's lively enough to stop your tot nodding off.

A handful of gentle Ogg soundtracks accompany the tasks, although there's little in the way of sound effects right now. Tasks are of varying complexity and arranged into categories – reading, maths, using the computer and others. Most are built-in to the main package, but others (such as the Chess game)

rely on external tools being installed. Thankfully, when something's missing, the program throws up a friendly error splash, so the child doesn't need to understand `errno.h` to dismiss it.

Learning zone

The reading tasks involve matching words to pictures, finding missing letters in words, listening to pronunciations of letters and various other relatively straightforward jobs which aren't too taxing. To give the child a sense of progression, most tasks offer a number of difficulty levels – in the picture matching, for instance, the easiest level throws up completely unrelated words, but on harder levels it uses words with similar spellings and/or meanings.

Tasks in the 'Discover the Computer' section educate the child on very elementary computer operation (moving the mouse, typing, drag-and-drop etc), and steers clear of advanced topics (few kids decode OOPS messages these days). In terms of maths, the Algebra task group includes working with money in Euros, constructing sums to match a result, all with basic addition, subtraction and multiplication covered.

Many of the games take a hands-on fun approach: helping a parachuting Tux land safely, piloting a submarine and escaping from a 3D maze all encourage a child to enhance his/her keyboard dexterity and mouse control. Moreover, the puzzles (memory, image matching and a simple little *Tower of Hanoi* variant) assist in developing a child's intuition and problem-solving abilities.

Apart from some wonky 'Franglais' in places, the only issue is that more feedback for successful actions could boost *GCompris's* presentation dramatically. Sounds or animations (or both!) would help to maintain a child's interest and provide an incentive to try other tasks. Our test infant didn't feel any particular thrill or satisfaction from completing a game – this needs to be worked on in future releases to motivate the end-user further. Otherwise, it's a remarkably well executed suite of happy little games and learning tools, and best of all... it's jam-packed with penguins!

EXPERIMENTAL GUI

SegusoLand

■ VERSION 0.11 ■ WEB <http://segusoland.sourceforge.net>

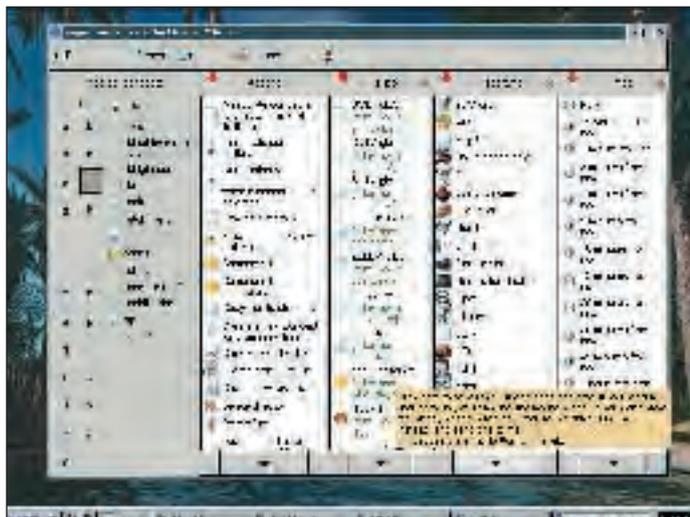
The X Window System, the basis for most Linux GUIs, was designed in such a way that window management policy and methods for starting applications were left to individual programs; as a result, today we have GNOME and KDE following many Windows-like GUI paradigms for newcomers, and different functionality for more adventurous projects such as *Fluxbox* and *Window Maker*. *SegusoLand* aims to make Linux far more accessible by associating actions with verbs and nouns (eg 'listen to music' rather than fiddling with menus, apps and commands).

To install *SegusoLand* you'll need a few extras installed beforehand – firstly, the *Boost C++* libraries (you don't need to compile, just specify where you've extracted it in the *Makefile*) together with *SWI-Prolog*.

Once they're both in place the compilation should go smoothly, and all the necessary files are available on our coverdisc. The documentation details it all thoroughly.

When first started, *SegusoLand* presents a fascinatingly constructed interface, with columns of files, programs, actions and times all bunched together. It's worlds apart from the traditional file management and program launching setups, and yet it all works curiously well. As an example, clicking **f** on our test system brings up a multitude of commands and filenames, and then clicking the **f1gp.gif** image allows us to select the action of **edit** with the program **gimp**. So far, so PDA...

Similarly, starting from scratch and clicking **edit** again brings up a whole host of editing actions (edit picture,



SegusoLand shows what it can do with just the word 'file'.

.doc, plain text file and so forth). *SegusoLand* does a fantastic job of keeping track of files and relevant software, and associating tools with goals the user wants to accomplish. Selecting 'listen' and then 'play audio' narrows the available files down to music, and considerate design touches like this are present throughout. *SegusoLand's* author has clearly put a

lot of thought into current GUI problems with Linux (and operating systems in general), and has also put forward a very strong case for including the utility in GNOME and KDE. Unquestionably, it makes first-time use of a new OS orders of magnitude simpler, and to experienced users it feels a more 'natural' way to operate a computer. Well worth investigating!

GRAPHICS PROGRAM

KolourPaint

■ VERSION 1.0 ■ WEB <http://kolourpaint.sourceforge.net>

Without a doubt, The GIMP is the best known graphics tool for Linux and was one of the OS's flagship apps in the early days. It straddles an unusual line between *Photoshop* and simple tools – not quite as featureful as the former, but requires more knowledge than the latter. We've had a friendly solution in the form of *XPaint* for many years, but its interface is somewhat archaic now and *KolourPaint* provides a fresh alternative.

KolourPaint strives to be 'conceptually simple and targeted towards the average user'. You'll need KDE 3.2 to compile it from source (along with an equally recent *Qt*), and it's built with the standard **./configure && make && make install** process. There's an RPM for SUSE on our coverdisc; you may find this works with other distros too.

Immediately noticeable is *KolourPaint's* aping of the *Windows Paint* interface layout – this is a good or bad thing depending on your point of view, but it's certainly wise of the coders to adopt a GUI that most will be familiar with. A bunch of brushes and shape tools are accessible down the left-hand side, while beneath sits the palette. *KolourPaint* can read a decent array of file types (JPG, BMP, GIF, ico, XPM, TIFF, EPS and others) and write to a similar range, although GIF saving is notable by its absence.

Alongside the basic drawing functions are additional features including crop, rotation, flip, skew and greyscale conversion – good enough for simple photo manipulation jobs; and the 10-level undo/redo and zoom facility are useful touches on top. There's no documentation to speak of at present, nor are there many



Familiar window furniture and an easy-going layout in KolourPaint.

configuration options, but the entire application is a cinch to operate.

All in all, *KolourPaint* doesn't come close to challenging *The GIMP's* mighty status in any way, but that's not the author's intention: as a simple,

easy and approachable paint tool it does a praiseworthy job. Here's hoping we'll see it rolled in to the next KDE release, as it's cleaner and healthier feature-wise than the venerable but decrepit *KPaint*.

SERVER MONITORING TOOL

HotSaNIC

■ VERSION 0.5.0-pre5 ■ WEB <http://hotsanic.sourceforge.net>

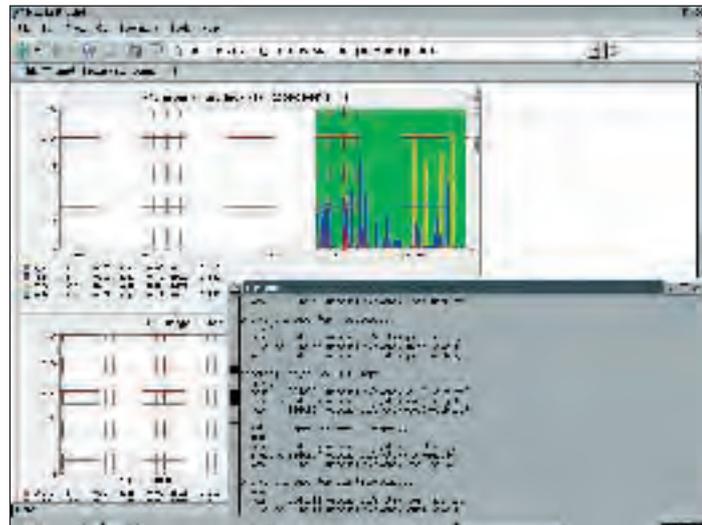
The most natural way to keep tabs on a server is by logging in to the machine with *SSH* (or, if desperate, good old *Telnet*), checking log files and poking around in */proc*. Sometimes this isn't the most appropriate method, though, and *HotSaNIC* (*HTML overview to System and Network Information Center*) provides an at-a-glance indication of a box's situation.

Installing *HotSaNIC* is fairly straightforward providing you have the *rrdtool* utility in place (and associated Perl libraries with **make site-perl-install**). From there, running **./setup.pl** will give a list of questions – after selecting which monitoring modules to enable, you can edit the 'settings' file in the main directory to specify how often images are generated etc. Lastly, running **./rrdgraph start**, waiting a few minutes for the databases to be

generated and then **./makeindex.pl** will produce the HTML output. All of this activity can be logged to spot any errors, if they occur.

In its default configuration, *HotSaNIC* can monitor network traffic over various interfaces, system load (memory, swap, disk, CPU and uptime etc), partitions and running processes, among others. SNMP support is included for monitoring traffic on remote machines, while scanning for worm hits on *Apache* logs and keeping track of a system's health with *lm_sensors* is also available. It's a solid range and covers most areas a typical admin wants to keep an eye on.

HotSaNIC's main HTML page provides thumbnail image links to full-sized graphs of the modules specified earlier; these are clear, colourful, well-rendered and suitably detailed. Graphs



HotSaNIC's CPU usage graph, with images being generated in foreground.

are produced for varying intervals depending on the module – the System section, for instance, shows info for the last hour, six hours, day, month and year. *HotSaNIC* is pretty speedy at generating the HTML and images, so it shouldn't impact performance on any moderately powerful server.

As a quick and simple Web-accessible method for monitoring a server's health and activity, *HotSaNIC* excels in providing frill-free layout and good images, allowing admins to spot problematic behaviour without hassle. If you're expecting a Slashdotting at some point in the future, prepare to be gazing at vast tracts of red...

AUTOMATED BATCH COMMANDS

FileGarden

■ VERSION 1.2.2C ■ WEB <http://redstone.sourceforge.net>

Those who've spent a lot of time at the UNIX command-line will appreciate just how versatile and powerful it can be. Being able to string together commands with pipes and redirection to make day-long operations take seconds is always rewarding, but to newcomers it all looks like a Pig-Latinised form of Martian gobbledygook. *FileGarden* attempts to make general batch commands much simpler, and is created with *Kylix* (the necessary runtime libraries are included on our coverdisc this issue).

FileGarden's install process is almost comically messy; even with many years of Linux experience under our belts and having spent countless hours installing random projects from the Net, we still couldn't get it running perfectly. The docs explain the essentials, but after

following them with pedantic precision we were *still* confronted by mystifying error dialogs, missing files and other

worries. It appears that the most salient problems can be resolved by installing and running as root, but if you find a fix, please share your knowledge with us by posting it on our forum at www.linuxformat.co.uk!

Installation woes aside, *FileGarden*'s smart GUI uses Red Hat's Bluecurve icons for added spit-shine and quick-access buttons down the left hand side for common tasks. It

doesn't fit in neatly with other desktop environments, but still avoids being too garish or unfathomable. *FileGarden* can spit out various bits of system information via the menu, and it's sane enough to filter such lengthy outputs as **last(1)** through **tail**.

By default, *FileGarden* is supplied with batch commands for converting images to different sizes and file formats with *ImageMagick* and *netpnm*, and writing text files as PostScript and PDF. Once the relevant files have been selected via the file browser, the quick access buttons or menu operations will perform the operation and display the output in the main pane.

Other operations can be added to the menus and buttons by editing *xfg.conf*; as with most aspects of *FileGarden*, your mileage will vary depending on installation, the runtime libraries and phase of the moon. *FileGarden* is an exceptionally promising tool which could help novices to grasp the command-line's power, but it needs some tidying up first. An intriguing project that we certainly haven't heard the last of...



FileGarden assisting in image conversion, with the Help box popped up.

COLLECTION MANAGER

Bookcase

■ VERSION 0.8 ■ WEB www.periapsis.org/bookcase/

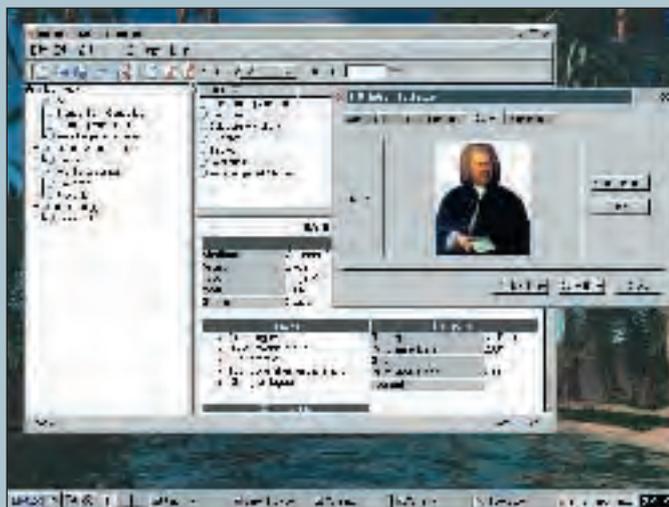
Unless you have a memory the size of Cheshire, managing a large collection of anything can turn into a brain-bending catastrophe. Trying to pinpoint that vital item can involve lengthy searches through disc inlays or book blurbs, so consequently tools for organising and searching collections can be a great help. *Bookcase* is one such utility, and a very fine one it is too.

We've provided an RPM of *Bookcase* on our coverdisc, created for Mandrake 9.2, but it may work on other recent distros including Fedora Core and SUSE. Failing that, compiling from source should be fairly straightforward providing you have the development headers for KDE and Qt (at least version 3.1 of each) along with *libxml2* and *libxslt1*.

When first started, *Bookcase* throws up the three-pane main window in which you can edit a

current collection or start a new one. For the latter, a satisfyingly thorough array of pre-set categories are available: books, comics, music, coins and even wine. The program's real versatility comes through with the custom collections, though – you can define all manner of fields, together with their types (text, lists, images and tables), which should cater for even the most bizarre collections.

The pre-set collections are remarkably well done, with a comprehensive range of fields to fill in and options to add cover images for books and music albums. Clearly a great deal of work has gone into supplying the user with everything needed; while the amount of fields can be overwhelming at first, the display pane renders selected items cleanly and coherently. Extensive searching facilities are included, and *Bookcase* can import data in BibTeX and CSV



Bookcase turns the chore of organising messy collections into a joy.

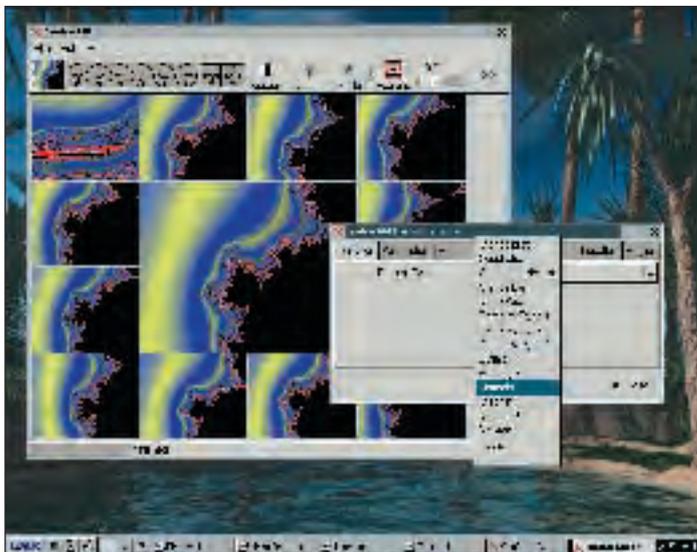
format – handy if you've already got a load of your info in a spreadsheet app.

For exporting, the program sports its own native Zipped XML format along with HTML, CSV and even the PalmOS database format. Despite the odd glitch here and there, *Bookcase* is maturing

well and is a first-rate method of organising collections of any kind. It has all the hallmarks of a great app: simple, flexible, fast and pleasant to work with. If – like us – you're having a nightmare keeping track of books, DVDs and CDs, it's a godsend!

FRACTAL IMAGE GENERATOR

Gnofract 4D

■ VERSION 1.9 ■ WEB <http://gnofract4d.sourceforge.net>

Gnofract 4D in Explorer mode, with the Settings box listing available types.

Fractal images are without doubt some of the most popular desktop backdrops, with incredible patterns and intense colour combinations far more lively than the old 'clouds.bmp' or 'xsetroot -solid grey'. Nutshellised, fractal images are generated from a complex number formula – each pixel is fed into the formula, which is then repeated over and over, and the eventual colour of the pixel is determined by how rapidly the result heads towards infinity. Stacks of fractal image generating software exists, but *Gnofract 4D* has a trick up its sleeve...

As it's built on GNOME you'll need all the usual development packages to compile *Gnofract4D* from source – in most mainstream distros these have names like *gnome-libs-devel* or similar. If you have no luck there, the RPM we've included on the cover CD should work providing all dependencies are resolved.

Gnofract 4D's *raison d'être* (and its most prominent feature over the original *Gnofract*) is its ability to produce images which are a cross between the Mandelbrot and Julia sets. As the author puts it: "it treats the sets

as different views of the same four-dimensional fractal object." Wowzers. Before your head explodes, we should point out that it isn't essential to understand all this mathematical babble – you can simply experiment and create fascinating images!

With lovely crisp rendering and support for 16 fractal types, *Gnofract 4D* is certainly more versatile than many similar tools we've seen, and it has a few other nifty additions too. The most notable of these is the Explorer mode; this slices up the main window with thumbnails of randomly altered 'mutant' versions of the central image. For finding the most impressive pictures, there's nothing better, and the abundance of alternative colour schemes is superb for getting the most out of an image too.

Gnofract 4D's options box sports a plethora of tweakable items, and these are all stored in the plain text param.fct file. Images themselves can be saved as PNGs, while the online documentation is sufficiently detailed for this kind of program. Great fun to just play around with.

MAIL SERVER VIRUS FILTERING

Sanitizer

■ VERSION 1.141 ■ WEB www.impsec.org/email-tools/procmail-security.html

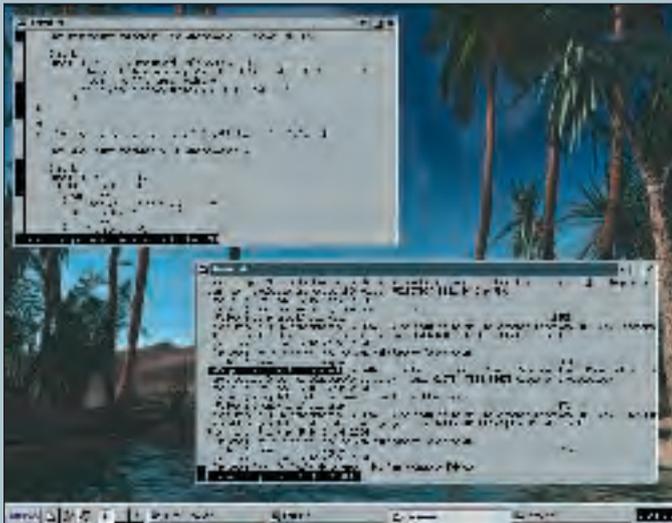
Spam and viruses-by-email have a lot in common: they're both unsolicited junk messages containing garbage, and they both cost companies vast amounts of money in wasted time. However, viruses are usually far more dangerous. When an inexperienced user blindly assumes that this attachment really is the latest Windows Service Pack, all hell can break loose. *Sanitizer*, then, is an email 'sanitizing' system designed to lessen the spread of viruses and worms.

At heart, *Sanitizer* is a *Procmail* script so you'll need that installed first (just about every server-oriented distro under the sun includes it). The tarball contains a number of scripts and HTML docs – see the configuration.html file for an example *procmailrc* with realistic settings. Information on integrating

with *Sendmail* and mail relays is provided in the well-written and marvellously detailed documentation.

With everything in place, and *Procmail* as the local MDA, the *Sanitizer* scans messages for suspicious content and tries to reduce their potential for damage. For example, an email containing 'active' HTML (ie scripts) will be 'defanged' and automatically tweaked, resulting in a less harmful version. These operations can be seen in the *procmail.log* or in messages themselves – an 'X-Security' header is added to the mail with the message 'sanitized on <host>', as is a link to information for confused recipients.

Sanitizer is capable of much more though, with the ability to strip executables and search inside Zip files topping the list of life-savers. In those emails where the executable's file type



Sanitizer's underlying workings and an example from the log up front.

has been obscured by an extra extension (eg *virus.exe.gif*), *Sanitizer* will 'mangle' the filename into a more recognisable state, and it can also check filenames against a text list of 'poisoned' files – easy to update when the latest virus appears.

As Linux viruses and worms are currently rarer than Horlicks Tablets, *Sanitizer* isn't really intended to be a

Linux end-user tool – it's primarily aimed at mail server admins with a load of Windows clients, and works crackingly well. And even though the script's author describes it as 'baroquely complex', it's fast and well-commented, making it easy to tweak week after week, or whenever a new wave of Windows email nastiness gallops into the wild...

CRONTAB EDITOR

GCrontab

■ VERSION 0.8.0 ■ WEB www.arquived.es/users/aldegado/proy/gcrontab/

Cron is one of those venerable Unix components that still plays an important role today. Unlike *atd(8)*, which does a single task at a specific time and date, *cron* runs regular background system tasks – every hour, day, month, week or any other set interval. On servers, these tasks traditionally involved compressing log files and building up the *locate(1)* database among other things; on modern desktops the tool is often used to rebuild font caches and prelink binaries. These tasks are defined in a *crontab* file.

GCrontab – suggested for inclusion here by LXF reader Marcus Jenkins – strives to make the editing of *cron* jobs easier via a slick GUI. It's built around version 1.2 of the *GTK* toolkit, as supplied with virtually every distro now, and has no other major dependencies.

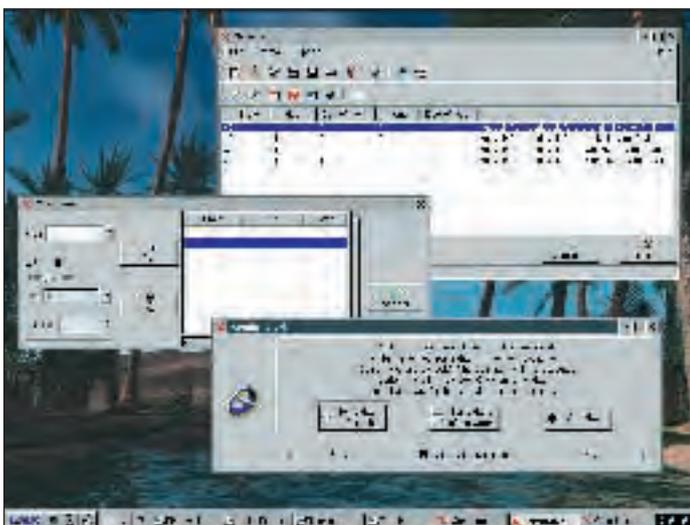
You may find an alarming amount of compile warnings with *GCC 3.x* – they're not fatal though, and won't affect the program's operation.

On most Linux installations, the user and root crontabs can be found in */var/spool/cron/crontabs* with appropriate read/write permissions. Running

```
gcrontab <file>
```

will display the main window; it's a compact setup sans pointless gadgets, and the tooltiped toolbars make it fairly easy to get started with. The current *crontab* is editable via a table, with the toolbar buttons providing quick access to modification operations.

Annoyingly, many of these dialogs don't present the user with the current values and data – instead, being totally empty – so you have to keep an eye



GCrontab busy at work, with the wizard-like *Druid* sitting at the bottom.

on things. Only a minor grumble though, as *GCrontab* redeems itself with its *Druid* feature – a wizard that steps through each phase of creating a *crontab* entry, with back/forward buttons and useful help text snippets. There's a sprinkling of other online documentation, giving an overview of the tool's basic workings.

GCronTab could do with a little more refinement in places, but it performs well and is no-nonsense in use. It's always refreshing to see smart little admin tools which don't depend on all the KDE or GNOME libraries, and *GCrontab* can quite realistically be installed on a server without munching up too much disk space.

CITY CONSTRUCTION

Senken

■ **VERSION** 0.3.0 ■ **WEB** www.contrib.andrew.cmu.edu/~tmartin/senken/

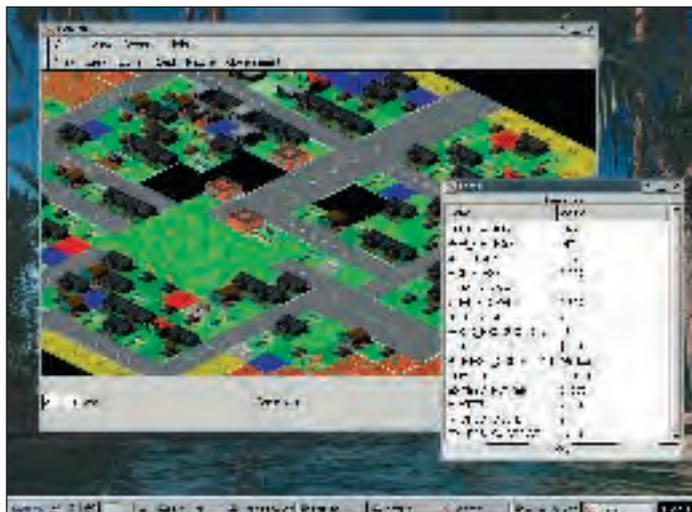
If you've ever spent seven hours in an unfamiliar place just searching for the Tourist Information Centre, you'll agree that some town planners are utterly *begging* for a visit from Mr Cluebat. Unfortunately, unless you're willing to devote some time dabbling in local politics, it's difficult to move all the best eateries and pubs onto your street (along with a few police stations to dampen the resulting chaos), but games like *Senken* give you free reign in a virtual world.

Senken is a town construction game (or 'mayor emulator') in the same vein as the ever-popular *Sim City* series *et al*, with all the usual chores and hassles the job entails – keeping employment up, crime down and population growth somewhere in the middle. It's a breeze to build from

source providing you have *SDL* and *GTK* installed (should be already included with most distros).

On first startup, *Senken* pops up a box to launch a game server (port 5000) to which you can connect the main game itself. Multiplayer mode is due to be added in future releases, but for now you're offered a choice of goals: grow cash reserves or population to a certain amount by a set date, or just start from scratch and build a town. These are similar to the Scenario modes in *Sim City* and provide some engaging challenges.

As expected, as the town planner you're responsible for laying out roads, building zones and supplementing them with services such as police and fire stations. Your goal is to keep citizens happy; their feelings and



A few more trees should, er, cut down the policing problem. Somehow.

problems can be gauged via the stats window, so if unemployment is high, you'll want to build more industrial and commercial zones, and so on. As with most games of this ilk, it's strangely compelling and a great endorphin-releaser when things are going well.

Senken's author evidently prefers a more gung-ho style of mayorship – the

citizens of your little community don't feel "content" or "dissatisfied"; but "cool with that", or "totally pi**ed!" (in the American sense of the word, obviously). Plenty of effort has been put into the game's workings and in general it's polished and enjoyable. Some features (*eg* saving) aren't fully implemented yet, but it definitely has great potential.

PLATFORM GAME

BlobWars

■ **VERSION** 0.4 ■ **WEB** www.parallelrealities.co.uk/blobWars.php

It's a hard life being a blob. Not only do you have no set shape or form, nor any ears to listen to The Brilliant Green, you're also under threat from sinister alien races.

BlobWars, from the talented folks at Parallel Realities (see *LXF39* for the slick shooter *Starfighter*) puts you in the shoes of Bob. Yes, Bob the Blob. As if having no arms was bad enough, Bob finds his world being invaded and decides to save the day with his heroic shenanigans.

As *BlobWars* is an *SDL* game, you should encounter no compilation niggles on any distro, although there are two specific dependencies you'll almost certainly need: the *SDL TTF* and *ZZIP* libraries. These are provided on our coverdisc, and if you build these first you'll have no problems later on.

We've come to expect lush presentation from Parallel Realities'

works, and *BlobWars* doesn't disappoint. The funny and to-the-point intro sets the scene for a great action romp, which features fitting background ditties and cleanly drawn sprites throughout. You're offered a number of difficulty levels and nine stages to progress through; these are all available from the start, but for the final release we think they should open up gradually (so that newcomers to the game don't accidentally dive in at the deep end).

Bandana-wearing Bob can enhance his weaponry by shooting baddies and snatching their power ups, giving him some truly nasty guns at his disposal, and in each stage your goal is to rescue the MIA (Missing In Action) Blobs, many of which are cowering away with entertaining thought bubbles. In a few areas the level design isn't totally inspired – thankfully though the collision



Hah! Ketchup! That'll teach him for dissin' my lack of body parts.

detection is spot-on and enemy deaths are accompanied by wonderful squelching sound effects.

BlobWars may only be at version 0.4 right now, and the developers are still fiddling with a few gameplay aspects, but already it's showing all

the signs of a Blobtastic platform shooter. It's polished, confidently executed and playable (Bob could do with some momentum though), and brings back warm memories of *Cauldron II* and the *Psion Jumpy* game. Excellent stuff. **LXF**

ULTIMATE LINUX BOX

Build the ULTIMATE Linux box

cover feature



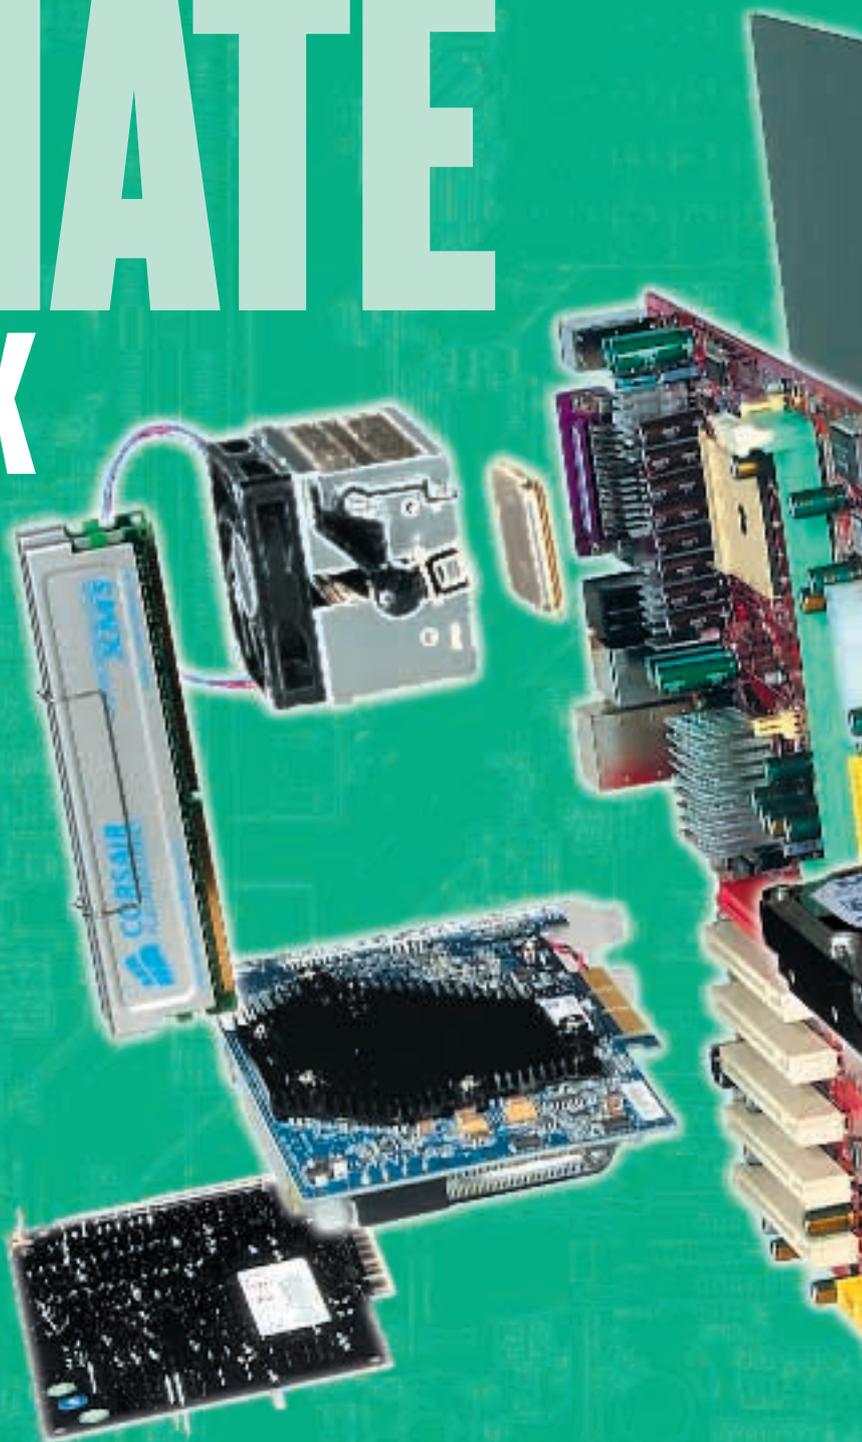
Ever believed the misconception that Linux doesn't play nicely with high-end graphics hardware? **Paul Hudson** sets about proving it's time that particular myth got blown out of the water...

An all-too-common scenario here at Future Publishing is some misinformed loudmouth from one of our sister gaming magazines exclaiming: "Linux and games? You'll be lucky!" and waddling off with a smug grin on their face. They'd say that Linux doesn't do decent graphics, that Linux can't do 7.1 surround sound, and that Linux and games was a nice idea, but we'd best just stick to *NetHack* and like it. Sure, some kudos to them for at least having heard of the venerable RPG that *NetHack* is, but on all other counts they couldn't be further from the truth, could they?

With our FUD-busting cannons primed with grapeshot and at the ready, the investigative journalists here at *Linux Format* took it upon ourselves to put together a machine that makes a first-class desktop platform for 2004: one that shows precisely what Linux is capable of across the board, leaving nothing to chance. This is the PC that most people dream about running a desktop system on – this is the Ultimate Linux PC for 2004. >>

IN ASSOCIATION WITH

SIMPLY
computers
www.simply.co.uk



ULTIMATE LINUX BOX



ULTIMATE LINUX BOX



SPEAKERS
Creative Labs T7700 7.1



GRAPHICS CARD Gigabyte GeForce FX5950 Ultra with 256MB DDR RAM



HARD DISK 2 x Maxtor
MaXLine Plus II 250GB



RAM 2 x Corsair XMS 512MB
DDR3200 RAM



HEAT SINK Cooler Master
Athlon 64 heatsink and fan combo

INVENTORY COMPONENTS

Every part hand-picked by the *LXF* team and the experts at Simply.

CPU

The defining part of any system is the CPU, and the choice here has a huge impact on choices made elsewhere. Although the Pentium 4 Extreme Edition, essentially a rebadged Xeon, makes for good performance, there was no doubt in our mind we wanted 64-bit capabilities. Although the Opteron is a great chip, this is a desktop machine, so we chose AMD's Athlon 64 3400+.

Yes, the FX chip range does provide a speed boost, but we're not sure it's worth the extra money unless you really need the higher performance. If you're a real extreme computing user, one benefit of the FX chip is that you can overclock it – though why you might want to is beyond us!

Plumping for 64-bits is important for two reasons. First, we know the Opteron is the highest-performing x86-compatible chip on the market, and we have high hopes for its younger sibling, the Athlon 64. Second, Linux has been the only x86-64-compatible OS available for over a year now, with no

64-bit release of Windows in sight. Although it's not something we'd usually do, just this once we want to thumb our noses in *PC Format's* direction!

OUR CHOICE: AMD Athlon 64 3400+

MOTHERBOARD

As the AMD64 architecture is still quite new, the mainboard selection is somewhat limited. Of course, that doesn't mean the options are poor – in fact, as the Athlon 64 has such an appeal to gamers who want maximum performance, all the boards come with a cracking selection of features.

Our choice – simply due to the amount of functionality packed onto it – was MSI's K8T Neo board, which has everything you could wish for and then some. Far from it from us to wax lyrical about a chunk of laminate, but this particular chunk comes with support for DDR400 RAM, AGP8x, Gigabit Ethernet LAN, 5.1 channel audio, eight USB ports (six at the rear, two on the front), as well as FireWire. Furthermore, it comes with a Promise Serial ATA RAID controller onboard

that supports RAID 0 or RAID 1, which should be able to squeeze the most performance out of this machine.

Finally, the board comes with MSI's CoreCell chip, which works with AMD's impressive Cool-'n'-Quiet technology to dynamically change the frequency of the CPU and speed of the fans based upon system use. If your machine is sitting idle, the CoreCell will underclock the CPU and slow the fans, thereby consuming less power and generating less heat. The second you start work again, the CPU gets brought back up to full-speed, and the fans returned to their usual state – CoreCell can even safely overclock the CPU by up to 10 per cent under heavy loads so that the machine really pulls its weight.

What other board can match this featureset at such a great price? No other boards, which is why the MSI K8T Neo was the clear winner.

OUR CHOICE: MSI K8T Neo

RAM

Although it's easy to go overboard on RAM, particularly with the 4GB limit

so handily smashed by the Athlon 64, we plumped for a relatively modest 1GB of DDR PC3200 RAM. Although RAM is a hotly contested marketplace, our choice had to either be Hyper-X from Kingston or XMS (Xtreme Memory Speed) from Corsair, both of which have an excellent pedigree and both can be relied upon in even the most demanding scenarios.

Of the two, we plumped for Corsair, and there are two key reasons for this. First, the XMS RAM comes with a platinum-coloured heat spreader unit on each RAM stick, which helps dissipate any excessive hot air. Secondly, XMS RAM has won more awards than we can shake a stick at, several of which come from sources we trust here at *Linux Format*.

PC3200 RAM (DDR400) is the fastest standardised RAM specification available right now, which should help keep the CPU well-fed with data.

OUR CHOICE: 2 x Corsair XMS 512MB DDR3200 RAM

GRAPHICS CARD

With the new range of graphics cards from nVIDIA and ATI due any day now, there are some real bargains to be had with their current models. There are a number of things to take into account

**MOTHERBOARD** MSI K8T Neo**CPU** AMD Athlon 64 3400+**SOUND CARD/
CONTROL BOX**
Creative Labs
SoundBlaster
Audigy 2 ZS
Platinum Pro**OPTICAL DRIVE** TDK Dual-format DVD-RW/+RW

here: 2D performance, OpenGL performance, price, and driver compatibility; of which the most important on Linux is the latter. With the majority of games still being for Windows, both nVIDIA and ATI have been competing heatedly for marketshare in that arena. On Linux, however, this is much less pronounced, which means drivers may be less used, less optimised, and updated less often than their Windows counterparts.

This is what makes driver compatibility the key to our choice, and in our experience we've had a much smoother installation experience when using nVIDIA cards, and this is what counts the most. The general consensus is that nVIDIA cards have the edge on ATI when it comes to OpenGL performance, and this, combined with the extra-easy driver installation, is what carries the card to success.

We anticipate that this will be the easiest part of the system to configure – partially because driver installation is straightforward, but also because the commercial version of Mandrake comes with the nVIDIA driver provided as standard.

OUR CHOICE: Gigabyte GeForce FX5950 Ultra 256MB DDR RAM

SOUND CARD

Too many people take the audio experience for granted. Sure, the bullets and helicopters aren't *real*, but that's no reason to pump their sound effects through the shoddy speakers that came free with your monitor. Our chosen monitor didn't come with speakers, but we were going to choose the best audio out there regardless.

Although one or two cards compete for the title, the undisputed champion in the area has long been Creative Labs with its Audigy line of sound cards. The latest in the range is the Audigy 2 ZS Platinum Pro model, which comes prepared for 7.1 surround sound at 24-bit/96KHz and plays back DVD-Audio to boot. One of the best little tweaks to this unit is the external unit that connects up to the PCI card through a special connector. This provides input and output ports for easy access, as well as a remote control sensor for the provided controller and even an extra FireWire port. The card itself has the usual speaker and headphone ports, as well as another FireWire port, which gives a huge amount of connectivity.

Of course, this unit is all about sound, and even though Creative doesn't make drivers for Linux yet, we

think we should be able to make this produce some top-quality audio with just a little effort.

OUR CHOICE: Creative Labs SoundBlaster Audigy 2 ZS Platinum Pro

HARD DISK

At long last, Serial ATA has settled itself into the hard drive market sufficiently enough for it to be called ubiquitous in new PCs. Although first-generation SATA only offers minor gains over standard ATA drives, it's still enough to make us choose SATA every time. Modern PCs released before SATA often used special 'rounded' IDE cables that squeeze the wide IDE cable down to a much smaller size, which, although a big improvement, is still about ten times larger than the micro-sized SATA cable. Therefore, from an aesthetics point of view as well, SATA wins out.

One final advantage to SATA is that the motherboard we chose comes with a Promise SATA RAID controller onboard, which means we can hook up RAID 0 or RAID 1 (striping without parity or mirroring) just by plugging the drives into the motherboard.

We chose two of Maxtor's new MaXLine Plus II 250GB SATA drives

for several reasons. First, at 250GB a pop, this machine has enough space to easily handle its data requirements for the next five years or so. If 500GB isn't enough for you, then you probably need to be shown how to delete files – all but prolific DV enthusiasts should find this sufficient! Second, the drives have sub-9.0ms access time, and an 8MB cache to keep the data transfer flowing smoothly. We expect optimum performance on all fronts with these drives, but if you want that little bit extra you can switch them over to RAID 0 (striped without parity) for that extra performance kick.

OUR CHOICE: 2 x Maxtor MaXLine Plus II 250GB

OPTICAL DRIVE

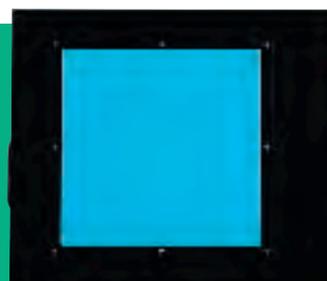
CD-rewriter drives are dirt-cheap in today's data-driven industry, which means that much of the pricing competition takes place in the DVD writer sphere. The big drawback here, of course, is that the DVD writing format has still yet to standardise on just one system, with DVD-R(W) and DVD+R(W) still fighting it out. Fortunately, increasing numbers of drive manufacturers are taking an ambivalent standpoint here and



ULTIMATE LINUX BOX



KEYBOARD Microsoft Wireless Optical Desktop



SIDE PANEL Essential? No, but it looks great! Thanks, Cooler Master!



MONITOR CTX PV722e 17" LCD



IN ASSOCIATION WITH
Simply
computers
www.simply.co.uk



TOWER CASE Cooler Master Praetorian

INVENTORY COMPONENTS

« adopting both formats, and this was what we wanted in our drive.

As the actual speed of writing is still quite slow right now, the biggest criteria – albeit rather superficially – is that the drive be black, and of a trusted brand. To this end we chose TDK's black DVD writer, which has 4xDVD+R, 2.4xDVD+RW, 4xDVD-R, 2xDVD-RW, 16xCD-R and 10xCD-RW. For reading, the specs are 12x for DVD-ROM and 40x for CD-ROM – more than enough.

By the time you read this, the all-new 16-speed dual-layer writers will just about be hitting the streets, which, thanks to the prevalence of dual-layer technology in purchased DVDs, essentially allows you to do a direct disc-to-disc copy for the vast majority of discs.

OUR CHOICE: TDK Dual-format DVD-RW/+RW

CASE

If you're going to spend £2000 on a PC, having it all wrapped up in a bland, beige box is a bit of a let-down. We wanted our machine to look the part, which meant we needed a case

that had an upbeat design whilst also insulating as much noise as possible. Although the market for cases is quite active, only one brand commands attention from everyone, and that's Cooler Master. We chose its Praetorian Black chassis for a variety of reasons: the 2mm-thick anodised aluminium material makes it strong, light, and attractive, and it comes fitted with four case fans as standard. Although we have no plans to overclock this machine – “really fast” is fast enough for us – having so many fans makes it a distinct possibility. Even the hard drives have two dedicated fans!

For CPU cooling, we used another Cooler Master product, its Athlon 64 heatsink and fan combo. If you're really into case modding, you would probably want to put some special thermal paste on the bottom of the heatsink, but we just used what was supplied. Again, we don't have any plans for overclocking! Though not the cheapest option, at £111, it's only a small fraction of the overall price and well worth it.

OUR CHOICE: Cooler Master Praetorian

MONITOR

This was quite a contentious category: partly because many people see monitor choice as an afterthought, and partly because of the fact that users' monitor choice is usually a subjective decision, and therefore very hard to pass judgement on with any degree of fairness. Having said that, we all spend most of our computing days staring at monitors, and really it's the most crucial choice out there: pick something poor and end up getting eye strain or headaches for years to come.

The choice between CRT and LCD was a no-brainer: although CRTs come in nice large sizes with prices to fit your pocket, they also take up entire desks by themselves and really don't look good in your home. Unless you have a cat that likes warming itself on the top of your CRT, LCDs are the way to go: slim, light, and much more attractive, they are worth the extra expense in our books.

We chose CTX's own PV722e 17-inch LCD screen for our computer primarily because it's light, has a first-class screen, and, on a side note, we wanted to see whether we could get the monitor pivoting to work with Linux – that's when you rotate the screen

90 degrees to one side so that it's in portrait mode rather than landscape.
OUR CHOICE: CTX PV722e 17" LCD

INPUT DEVICE

Back in *LXF46* we reviewed a selection of keyboards and mice to see how compatible they were with Linux. However, things have moved on since then, with the two key innovations in the market both coming from Microsoft. Yes, we realise there's an inherent dislike of all things Microsoft amongst some Linux people, but we're hoping you can put that behind and concentrate on the hardware here!

The first interesting innovation is the longer battery life. By turning the transmitters off faster and more frequently, Microsoft claims that the new keyboard and mouse will last up to six months on two AA batteries. The second innovation is the side-scrolling mouse wheel – something we'd certainly like to see working under Linux. As this is Microsoft hardware, it's almost certain to have no support as yet, and so it's pretty long odds on getting this to work – but we appreciate the challenge!

OUR CHOICE: Microsoft Wireless Optical Desktop

START THE BUILD

Building the Ultimate Linux box isn't ultimately difficult.

Consumers have always been lead to believe that computers run using black magic and incantations – after all, if you dare to open your PC case, you *void the warranty*. However, the reality is that computers are actually quite tough, and as long as you don't plan to carry components around whilst wearing socks and walking on a fluffy carpet, you should be fairly safe from static. Having said that, we are working with expensive kit, and it's better to pay £5 for an anti-static device rather than £500 for a new graphics card.

Most motherboards and cases come with comprehensive manuals detailing the workings of every pin and jumper – keep this to hand, and try not to lose it in future. At the very least, keeping the motherboard manual to hand is helpful if you make a fatal change to your BIOS and want to know how to reset the CMOS data. We strongly recommend you take the time to read through the manual – even though it is usually in-line with *LXF's Flex & Bison* tutorial on the Interesting scale – to ensure you don't

scupper your build project by making any elementary blunders.

Before we begin, one last note: don't be afraid to give things a fairly strong push/pull in order to get them to fit, but be warned that there are often small

hooks that need to be undone to allow safe removal of parts. The graphics card, for example, features a tough retaining mechanism that locks it into place safely – don't try to remove the card without first unhooking it!

“We strongly recommend that you take time to read through the manual to ensure you don't make any elementary blunders.”

STEP 01



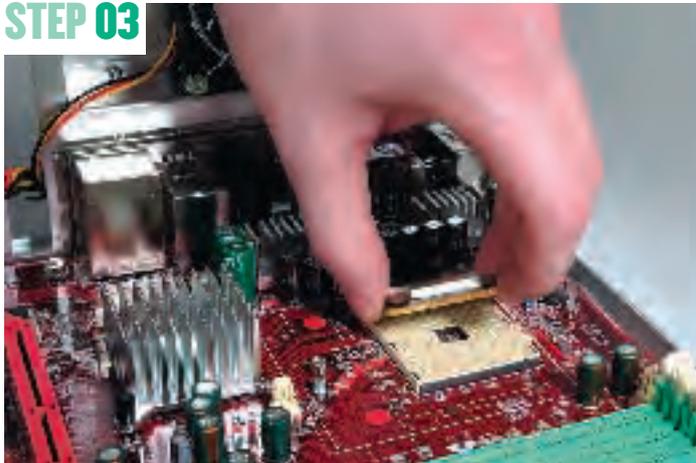
STRIP THE CASE In order to get started fitting the core components of the machine, you need to take both sides off the case, remove the chassis backplane, and untie all the wires. The Praetorian box has a handy caddie for the backplane and thumbscrews where possible, which makes all this quite easy, and screwdriver-free too.

STEP 02



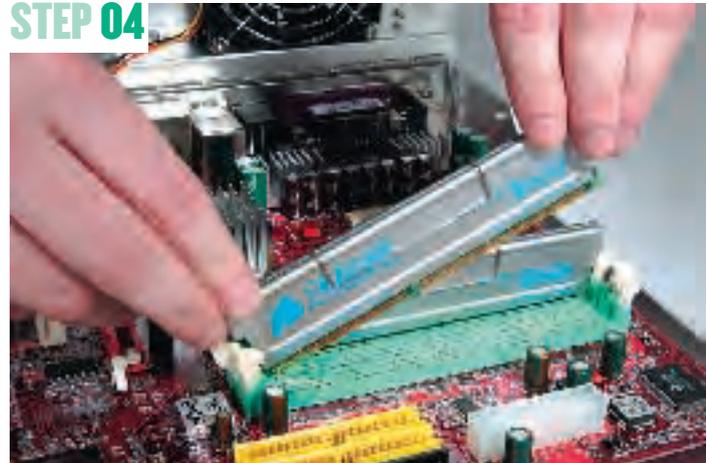
ATTACH THE MAINBOARD Wearing your anti-static wristband, screw the motherboard onto the chassis backplane. You'll usually need to put spacers into the board first and screw the motherboard onto them. The more screws you can use, the better – some cards need substantial pressure to fit. NOTE: the Athlon 64 CPU requires a heatsink backplate attached to the rear of the motherboard.

STEP 03



SLOT IN CPU As the pins on the CPU are often symmetrical, you should rely on a marking to tell you the correct orientation (in the Athlon 64's case, it's a golden arrow). Drop the CPU lightly into the ZIF socket, letting gravity help the pins fall into their holes, then lower the lever halfway down. Finally, press lightly on top of the CPU, then push the lever the rest of the way.

STEP 04

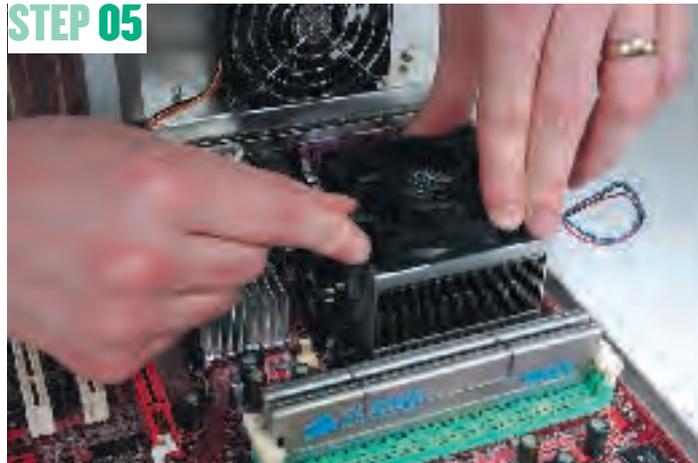


FIT RAM RAM sticks, like all smart PC components, are designed to only go in one way, so you should have no difficulty here. Our RAM sticks have heat-dissipation covers which also makes them safer to handle – unprotected RAM tends to be prone to getting static-damaged very easily. Fit your RAM in the first two free slots, and you're done.



ULTIMATE LINUX BOX

STEP 05



ATTACH FAN AND HEATSINK This can be quite difficult; or, if you failed to fit the heatsink backplate back in step 2, you will find that it is flat-out impossible! In our case, we attached a special plastic retention mechanism to the heatsink backplate, then clipped the heatsink on top. This contact needs to be very tight in order to allow heat transfer to take place properly.

STEP 06



REPLACE THE BACKPLANE CHASSIS With the CPU and heatsink fitted, it's now time to re-insert the backplane chassis so that we can start putting the rest of the components in. As experienced case-modders and PC builders will know, you can – technically – do a lot more with the chassis out, but our fitting environment was far from dust-free, so it's just not worth the risk.

STEP 07



SCREW SIDE PANEL ON Again, as we had significant amounts of dust in our fitting environment, we decided to put the left-hand case panel on to keep the motherboard protected – unless you're super-scrupulous about cleaning your office/workplace, you should consider this too. We'll be removing the left-hand panel briefly later on to screw in drives.

STEP 08



FIT GRAPHICS CARD Like most other 5900/5950 cards, our GeForce uses two PCI slots because it runs hot enough that a big heatsink and fan is required to keep it cool. Using the extra slot isn't a problem – our motherboard has so much built-in that we have no shortage of PCI slots. It does need a dedicated power supply from the PSU to get anything beyond VESA graphics though.

STEP 09



SLOT IN SOUND CARD Pick a free PCI slot and use it for the sound card. NOTE that we put our sound card pretty much as far away from the graphics card as it could go, to give the GeForce card as much room to cool as it can. As there are quite a few wires to connect into the card if you want to (such as the optical drive), you might want to leave it at least one up from the bottom for easy access.

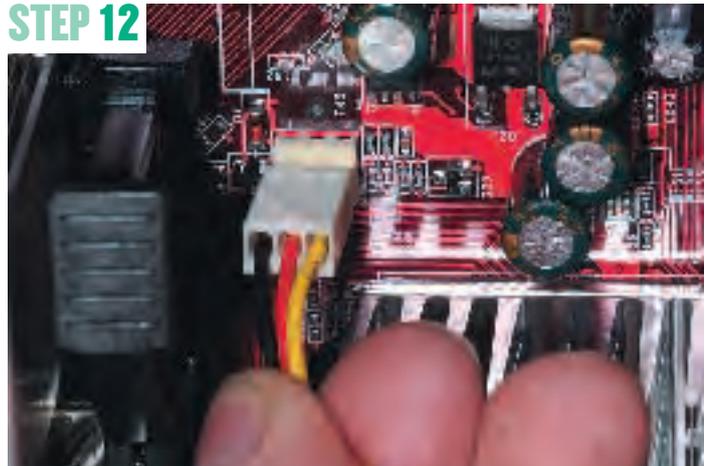
STEP 10



ATTACH HARD DRIVES Our Ultimate Linux box boasts two hard drives, and we fitted them quite low in the chassis. The reason for this is because the Cooler Master case that we chose has fans specifically for the hard drives, and we wanted them to get maximum usage. The drives were connected up to the Promise SATA RAID controller for maximum performance.

STEP 11

SLIDE IN OPTICAL DRIVE Where you place this is likely to depend somewhat on whether your machine sits on your desk or beneath it. However, don't forget to take into account the length of your IDE cable – if you only have something short, putting the optical drive in the top bay isn't a smart move!

STEP 12

PLUG IN CABLES Above, you can see us plugging in the main CPU fan to the motherboard, but we also plugged in the top, back, and front case fans (four in total), the CPU, graphics card, hard disk, and optical drive power, the CD-to-sound card cable, as well as the front ports. As there are at least a dozen cables to plug in at this point, we suggest you study your mainboard documentation carefully!

STEP 13

FIT FANCY SIDE PANEL This step is, of course, entirely optional, but we couldn't resist – our Ultimate Linux box includes a transparent blue acrylic side panel to show off its contents to the outside world. If you want any credibility in case-modding conventions or LAN parties, you'll need one of these. Our panel was provided by Cooler Master for this unit, and so fitted the case perfectly.

STEP 14

BASK IN THE GLORY The main machine is built, so whip out your monitor, speakers, and input devices, and finish your setup by stringing all the peripheral spaghetti to its respective ports. Unless you plan to use your computer inside a broom cupboard, you may wish to be slightly more creative (excuse the pun) with your speaker layout than we've been in the photo!

THE BUILD CONCLUSION

Now, that wasn't hard at all, was it?

If you haven't put together your own machine in a while, you should have noticed at least two differences heralded by the times.

Firstly, both the graphics card and CPU require their own dedicated power from the PSU. This should immediately give you a good idea of how much wattage they must draw, and thus how powerful they are.

Secondly, the heatsink and fan combo requires a special backplate to be fitted to the motherboard to hold it on. This never used to be the case,

simply because in less-demanding heritage boxes, the heatsinks used to be a whole lot smaller! Again, this update is simply down to the sheer power of the CPUs used in modern machines: as heat rises, so the need to get it dissipated rises too.

“Machine doesn't boot up? Don't give up hope – our motherboard comes with a special D-bracket for problem diagnosis.”

If your machine doesn't boot up, don't give up hope: it may not be a busted component. Our motherboard – remember, we did choose it because of the features it provided – comes with a special D-Bracket that can be fitted to the motherboard to

provide detailed information about the boot-up process. For example, if the CPU isn't fitted correctly, is broken, or is without its dedicated power lead, the D-Bracket will show four red lights. With sixteen different combinations of red and green lights available, you can instantly diagnose a problem and replace any poorly fitted or malfunctioning hardware.

If you've made it this far without any problems, it's time to get the OS installed and configured, which is where the *real* work often begins. >>



SYSTEM CONFIGURATION

Hardware tweaking and dual-booting into both 32-bit and 64-bit mode.

No matter how good your hardware is, if you don't take the time to properly configure the software

you'll get poor performance. In this section we'll be looking at how the OS was installed, how the graphics and

sound cards were configured, and finally how we tweaked the peripherals to get all the extra keys to work.



INSTALLING MANDRAKE

One of the main reasons we like Mandrake so much is that it has such an easy install procedure. With MDK 10 out now (and on this month's coverdiscs, too!) for 32-bit systems, we sorely wanted to use it on our Ultimate PC.

However, as we wanted maximum performance, we opted for the 64-bit release of Mandrake 9.2. This is essentially an updated 32-bit 9.2, with most of the bug fixes and updates pre-installed. Sadly it doesn't have KDE 3.2 or kernel 2.6 – we'll have to wait till May for the final Mandrake 10

AMD64 release to get these features. In the meantime, as we had two drives, we ended up dual-booting both 9.2/AMD64 and 10/x86, which gave us the opportunity to directly test how well the chip does in 32-bit mode compared to 64-bit.

Both install systems are almost identical, and both times we did a fairly complete install in order to avoid having to install things one-by-one later. This is, of course, one of the big advantages to having huge hard disks – 'minimal' installs are a thing of the past.

Mandrake detected all the hardware correctly, though it managed to get much higher resolution by default in Mandrake 10.0 due to the improved nVIDIA drivers in there – our card was released after 9.2, which makes support for it difficult. Similarly, the sound card was also detected, although the extra buttons on the keyboard and mouse were left untouched. We had one or two problems in the 64-bit 9.2 VIA SATA controller, but this wasn't a problem as we had our drives hooked up to the Promise RAID controller, which was

supported. As we were installing this, the first beta of 10/AMD64 was released, with the main new feature being improved support for SATA controllers from VIA and Promise. Figures!

Once everything was installed, the system booted up with no problems. To ensure we had no problems, we downloaded the latest BIOS patch for the DVD drive, only to find it was a Windows executable. Thankfully, Wine continues to prove that it's as good as we all thought it was, and the firmware patch was applied without hitch.

CONFIGURING THE GRAPHICS CARD

As mentioned earlier in the components section, we expected this part of the process to be the easiest – we weren't disappointed! First, go to *Mandrake Control Center* and install the *kernel-source* package. The nVIDIA driver itself is closed-source, with what is essentially a stub file left open. This stub file is the communication layer between the kernel and the nVIDIA driver, and so the nVIDIA installer needs to compile that stub file against your current kernel.

The first step in this process is to go to www.nvidia.com and grab the

latest drivers for your system. In our case, this was Linux AMD64.

Next, you need to close any open XFree86 sessions you have, switch to root, and type **sh ./<name of installer>** and hit **Enter**. This will launch the nVIDIA installer. Chances are it will say it has no support for your kernel, and will offer to search the nVIDIA FTP site. This is probably pointless, as it's not likely there will be support for your kernel on their FTP site either. Happily, this is not necessary – as we have our own kernel source,

the nVIDIA installer will compile and build a custom module for our kernel.

This should happen automatically, and when the installer finishes it should flash up a message about you needing to change your X11 config file. Type **vim /etc/X11/XF86Config-4**

and search for your current video card. It should look something like this:

```
Section "Device"
Identifier "device1"
VendorName "NVIDIA Corp"
BoardName "NVIDIA GeForce FX
```

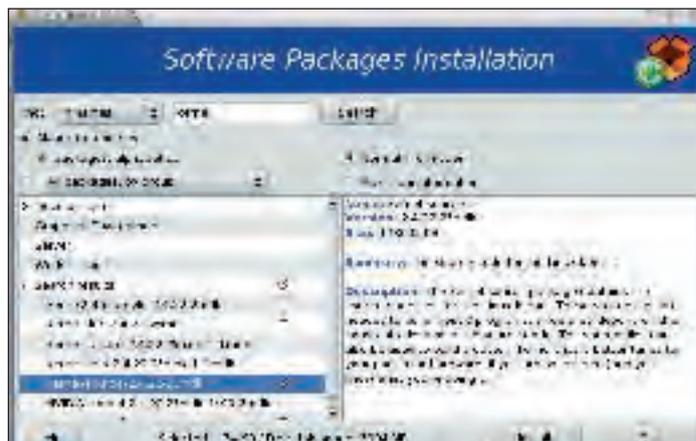
5950 Ultra"

Driver "vesa"

Option "DPMS"

EndSection

You'll need to change **vesa** to **nvidia** in order to activate the nVIDIA driver. Also, if you see any lines that load **dri** or **GLcore**, you should put a hash (**#**) at the start of the line to comment it out. Once you've saved that, you should be able to type **startx** to start back into X Windows with the nVIDIA driver enabled.



You need to install the source code for your kernel in order to make the nVIDIA driver work. This can all be done through *Mandrake Control Center*.



Whether it's 32-bit or 64-bit Linux, nVIDIA has a driver already made for you – this is the kind of Linux support that we like to see!

CONFIGURING THE SOUND CARD

This was one thing that required no effort on our behalf. Under 9.2, the card was detected by Mandrake and the *Emu10k1* driver was installed just

fine. Under 10.0, ALSA had great support for the card, which made configuration a non-issue. Games such as *Unreal Tournament 2004* and

Tux Racer worked first time with sound, both of which use SDL.

It's worth noting that if you do experience problems with any Creative

product on Linux, it's worth checking out the links at Creative's Open Source support page at <http://opensource.creative.com/>

PERIPHERALS

As always, Hoyt Duff's tutorial from *LXF28* was first to hand for getting all the niggly bits of our keyboard working. However, for the purposes of this feature, we'll cover how to get the Calculator button to launch *KCalc* when pressed. First, make sure you have installed the X11R6-contribs package, as this contains the *xev* package that acts as a keycode scanner.

The first step is to use *xev* to find out the keycodes of your special keys. Crank it up in a console, highlight the new *xev* window, and press a special key – you should see a keydown notification and a keyrelease notification something like this:

```
KeyRelease event, serial 27,
synthetic NO, window 0x2000001,
root 0x3a, subw 0x0, time
2573184, (85,68), root:(702,452),
state 0x0, keycode 228 (keysym
Ox0, NoSymbol), same_screen YES,
XLookupString gives 0 bytes: ""
```

What we're interested in is on the third line: keycode 228. That's the unique key identifier generated by this keypress, and we need to assign that to a useable key. This is done back in the console – from your home directory, type this:

```
cd .kde
cd Autostart
vim keycodes
```

In this file, add all the keycodes you got from *xev* along with the key you'd like to assign them to. I used F keys above F20 as they are definitely free, so here were the first three lines from my keycodes file:

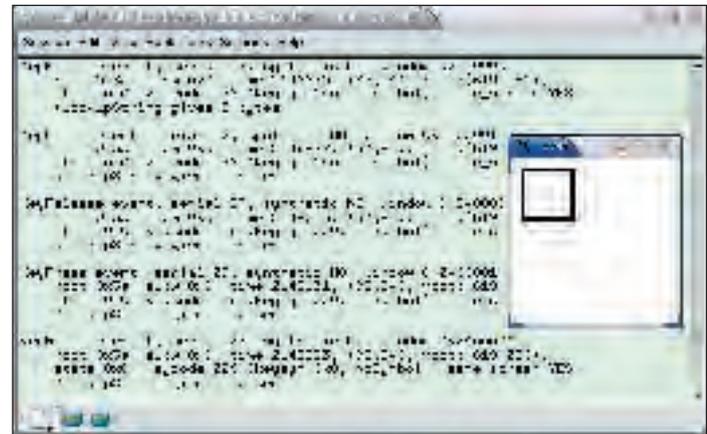
```
#!/bin/sh
xmodmap -e 'keycode 161=F21'
#calculator
xmodmap -e 'keycode 101=F22'
#My Documents
```

Once you're done – note that I've added comments after each key to remember what the key actually represents on the keyboard – save the file, then quit *vim* and type this:

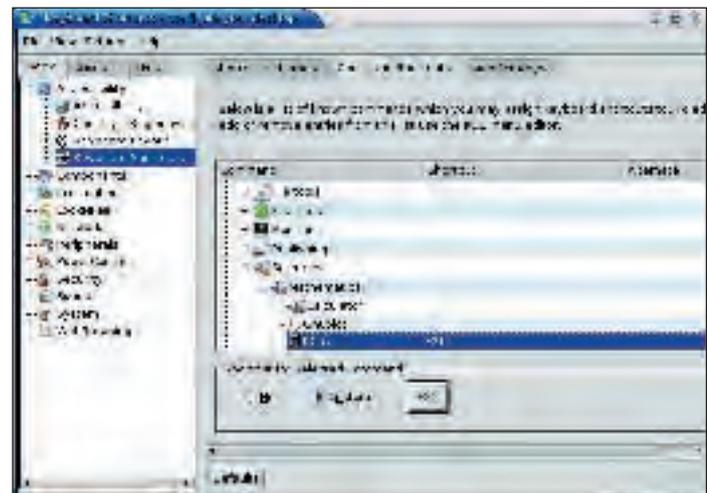
```
chmod +x keycodes
./keycodes
```

That will make *keycodes* executable, then run it. Every time KDE runs in the future, this will also be executed as it's in the KDE Autostart directory. Having run that file, your function keys should now be enabled, so we're on to step 2: making the keys actually work in KDE.

Thankfully, it is quite straightforward. You need to go to the *KDE Control Center*, select Accessibility, then Keyboard Shortcuts. From here you can make your new keys do pretty much whatever you like, but our plan is to make the calculator button (F21 in my file) load *KCalc*. To do this you need to go to the Command Shortcuts



When *xev* is running, it automatically outputs all X events (hence the name) to the console that launched it. This is the easiest way to track down all the odd scancodes generated by modern keyboards.



Assigning your new key to any program is easy, thanks to KDE's flexible and easy-to-configure hotkey system.



tab, look for *KCalc* in the menu list, then select it and choose the 'Custom' radio button from the options below – we want to assign a custom key for it, see?

When you hit Custom, a new window

should pop up asking you for the primary and second shortcuts. Unless you're really into control, you probably only need a primary shortcut – just hit the Calculator button on your keyboard and the screen should immediately go away. This is because the Calculator key has been mapped to F21, so you should see the previous screen updated to reflect F21 as the shortcut for *KCalc*. Hit Apply from the *KDE Control Center* dialog window, then try your key out!



« BREAKING IT DOWN

Before we even play our first game, what has the project taught us?

At this point we're done. The machine is built, the software is installed and configured, and we can at last sit back and enjoy the fruits of our labour. As you've seen over the last few pages, even getting the most advanced graphics and sound cards to work is just a matter of following a few simple instructions, so all the myths about Linux being hard to use are just that: myths.

Perhaps the hardest part of the system was getting all the keyboard buttons to work, but even that took less than five minutes, and you've got to admit it's pretty cool to be able to map any key to any function, so it was well worth the effort.

After having fitted everything and tested it all out, the question "What would we change if we had the chance?" springs to mind. Well, if money had been absolutely no object, we maybe would have gone for AMD's FX-53 chip. It's clocked slightly higher, and has a few architectural tweaks to make it go that extra distance for maximum performance. However, it's also about three times the price, so the hackneyed maxim "you get what you pay for" is one that seems to suffer from diminishing returns in this case. Similarly, we could have opted for the newer 10K RPM SATA drives instead of our chosen units, but these inevitably weigh in at higher prices too, and usually come with much reduced capacity. If we were super-hardcore, we might have had two 10K RPM drives striped together for the operating system and the two other drives striped together for data, but again: if you have to ask how much it costs, you can't afford it!

Lessons learnt

There are three key points that we thought were especially highlighted in the production of this feature:

- Although the CPU, graphics card, and sound card are important to system performance, the motherboard is what really has the biggest impact on how usable the system is. We chose the most feature-packed board, and it really made a big difference.

SUGAR AND SPICE AND ALL THINGS NICE

The components we put into the Ultimate Linux box for 2004, and how you can get your mitts on them yourself...

We chose all our components direct from the Simply Computers shelf, so here's the full list of everything that went in, along with the component number in the catalogue and the price excluding VAT.

DESCRIPTION	PART NO.	PRICE (EX.VAT)
AMD Athlon 64 3400+	BN75003	£247
MSI K8T Neo motherboard	BN49667	£82
512GB PC3200 Corsair RAM stick	BN64979	£51.99
512GB PC3200 Corsair RAM stick	BN64979	£51.99
Maxtor 250GB SATA drive	BN75109	£131
Maxtor 250GB SATA drive	BN75109	£131
TDK multi-format DVD-RW/+RW optical drive	BN49249	£87.99
Gigabyte GeForce FX 5950 Ultra	BN48993	£275.00
SoundBlaster Audigy 2 ZS	BN49684	£135.17
Creative Labs T7700 7.1 speakers	BN49679	£79.99
Cooler Master Praetorian case	BN49850	£111.00
Cooler Master heatsink and fan	BN5P26310	£8.00
Microsoft Wireless Optical Desktop	BN66822	£42.54
CTX PV722E 17inch LCD monitor	BN65284	£308.24
TOTAL		£1516.60

- Supporting manufacturers that put a lot of effort behind their Linux drivers, such as nVIDIA, helps the community, and you also get a more stable system as well.

- In a straight comparison between 32-bit and 64-bit distros, the 64-bit version will trounce the 32-bit version with no tweaking required. Having said that, we're still looking forward to seeing Mandrake 10 for AMD64 and all the kernel 2.6 goodness it brings with it.

The Ultimate in 2005 and beyond?

With the system complete, all that's left for us to do is consider what the ultimate machine will be this time next year, and, of course, to actually review the machine as if it were built and sent to us like it is. As for the ultimate in 2005, we have a strong feeling that it will continue to run on an Athlon 64 chip with nVIDIA graphics and Creative Labs sound, although things will certainly have moved on a little by then. The CPU is likely to be a minimum of a 4000+ (perhaps even a 4400+), and graphics card is likely to be a more souped-up version of the NV40 card. NV40 itself, due to be

launched by the time you read this, is likely to start at 'GeForce 6800' in the same way that the GeForce FX started at 5800, so perhaps this time next year the GeForce 6950 Ultra will be what's wowing high-end gamers.

We fully expect Intel to have a competing 64-bit desktop chip by this time next year, although it's anybody's guess how well it will perform against the Athlon 64. The AMD chip is a much more mature proposition, though, and that's likely to count heavily in its favour.

RAM is of course also likely to continue clocking up, and this time next year we expect to see 533MHz RAM as the emerging standard, with an even faster standard already on the horizon for higher-end systems.

Hard drives? Well, the 10K revolution (excuse the pun) should eventually reach drives with large capacities, so we'd like to see the equivalent machine next year with the same size of drives just with a much faster RPM.

For now, though, and also for the foreseeable future, our Ultimate box is all set to leave the competition choking on its dust. But just how well would our lovingly constructed chunk of hardware perform if it was to be subjected to the rigours of a normal LXF review, with its attendant benchmarking with our all-new LXF Bench 2004? Read on to find out!

THANK YOU

Simply Computers provided all the parts necessary for our build, and deserve a big round of applause for doing so much to help the Linux community. Both Sammy Laing and Atif Malik at Simply did a sterling job helping us choose the best-value components, and this feature would simply not have been possible without all their knowledge and hard work – thanks!





BENCHMARKING

How our Ultimate Linux box fares in a conventional *LXF* review...

There's no doubt it's hard to objectively review a machine you designed, tested, and configured yourself, but it just wouldn't be right for us to go through all this effort without telling you how the final machine performed!

Although we've already looked at each individual component in isolation, it's time to look at them as a discrete working computer. The Athlon 64 3400+ CPU is fantastic value for money, and, combined with Corsair's XMS RAM we're expecting some first-class benchmark results. Both the drives are SATA and hooked up to the Promise RAID controller, and so should be able to work at their peak capacity. At the time of writing, Mandrake 10 beta 1 for AMD64 just poked its nose out of the door, and a lot of work has been done to improve SATA RAID support. As we were using 9.2, we opted not to stripe the two drives together so as not to push our luck!

For graphics and sound, there really isn't much to say: there are simply no better choices on the market right now than the two cards we chose here, as both NVIDIA and Creative have managed to cram so much innovation and value into their products that they can't be anything but stunning. Increasing numbers of Linux workstations come with nVIDIA graphics cards, but few come with such meaty sound cards, and none come with our wall-shaking speakers! We're expecting to be blown away on the gaming front.

Despite having six fans (seven if you include the oversized fan on the GeForce card) in the box, this machine is only barely audible when turned on and running idle. This is largely thanks to the CoreCell chip on the MSI board, and you'll find that when under load you'll be able to hear at least a faint hum. Both hard drives run silently unless you sit with your ear directly against the case.

On a side note, back in *LXF46* the Microsoft keyboard/mouse combo we reviewed seemed to have a fairly poor transceiver – we had to place it fairly close to where the mouse and keyboard were to get flawless access, which almost defied the point of wireless. Thankfully, the glitch has been fixed in this newer model.

Making it zip

As we dual-booted Mandrake 10 for x86 and Mandrake 9.2 for AMD64, we had the opportunity to run our benchmarks in both 64-bit and 32-bit mode. This is something we've not yet had the luxury of doing, so we

were waiting in suspense for the final scores to turn up. Firstly, the hard drive score was 1.31, putting it at 33 per cent faster than our benchmark machine. The drive score varied only slightly between the operating systems, as you can imagine.

Now, onto the three CPU-reliant scores. Before we proceed, let us remind you that Mandrake 10 does come with kernel 2.6, which benefits from substantially improved performance as compared to 2.4. Now, in 32-bit mode this machine returned a RAM score of 2.44, a single CPU score of 2.15, and a multi-CPU score of 2.1. These are very respectable results, given that our yardstick machine is a 1.8GHz Pentium 4. However, as we've said numerous times, we expect AMD64 machines to perform more than twice as fast as their equivalently clocked Pentium chips, so a result of 2.1 is certainly a little below par.

Of course, this was all in 32-bit mode – booting up Mandrake 9.2/AMD64 was a different story entirely. This time the machine scored 3.15 for the single CPU test and 3.01 for the multi-CPU test, which is about 10 per cent above what we have come to expect from our previous tests, and equivalent to a 5.67GHz Pentium 4. Yes, this sounds a little crazy, but there's a reason the Athlon 64 performs better. Remember, all our previous tests were done on servers (with the odd workstation), all of which used ECC-registered RAM

with parity that certainly wasn't PC3200-spec. RAM performance makes a huge difference to overall system speed, particularly as the Athlon 64 doesn't need to bother with parity data, and is almost certainly the cause for this jump in performance.

Are these scores surprising? Not at all. While they are certainly above what you might expect, it's important to remember that we chose the best components across the board: it's a first-class motherboard and RAM combo, with the very latest hard drives from Maxtor, and AMD's flagship desktop chip.

Is it the ultimate?

After having slaved away for so many days to get everything tuned for performance and get every last bit of driver functionality to work, we can safely declare that *everything in this machine can and does work perfectly under Linux*.

Now, the final question to answer is: is this machine really The Ultimate for 2004? Last issue, that Systemax Mission 3602 server won a *Top Stuff* award because of its excellent performance and value, and it actually scored higher on the benchmarks than our Ultimate box. However, it's important to look at the benchmark breakdown – our Ultimate box has it beaten in the RAM tests by about a third, and is about three times faster for hard drive usage. The primary reason our Ultimate box didn't win on the overall score is because the Systemax box scored 6.07 on the multi-CPU test – it does, after all, have two Opteron chips in there. If you consider the Systemax 3602 with just one CPU, our Ultimate box roars ahead to be a comfortable winner.

So, in short, we fully expect this Ultimate machine to be the average machine in about two years time, and still able to play the newest games in three or even four years time – it's that good. [LXF](#)

NO OVERCLOCKING

Note that for the purposes of our review benchmark, the safe overclocking functionality of the motherboard was disabled. However, we did enable it at a later date to ensure it was indeed wholly safe, and experienced no problems irrespective of how much stress we put the system under.

BENCHMARKS – 64-bit

CPU	3.01
SINGLE	3.15
RAM	2.75
HD	1.31



What on Earth is... MONO?

The .NET development framework is one of the newest and more revolutionary ideas on software development and deployment. Why should we want the .NET framework ported to Linux? **Biagio Lucini** investigates.

» Mono? Doesn't that feel like the flu and make your glands swell up?

Mono is a project aimed at providing an Open Source implementation of the Microsoft .NET development framework.

» Good, but what is Microsoft .NET?

What is called with generic terms "Microsoft .NET" is in fact the union of different projects and strategies from Microsoft. In particular, we have:

- a** development framework
- b** Passport and other services
- c** software and products with the .NET brand

The Mono effort targets just the development framework, *ie* an ensemble of tools for developers, libraries and a runtime environment conforming to the .NET specifications. To avoid confusion, let's say it explicitly: Mono *won't* provide us with an implementation of Microsoft Passport or other .NET services, *nor* will it allow the Microsoft .NET branded products (like Office .NET or VisualStudio .NET) run on Linux; if you are interested in this, you should look elsewhere. However, if you are a developer and want to use the .NET framework, or you are interested in running a .NET application written with this framework, Mono is the answer.

» Is Microsoft helping in the development of Mono?

There is no official cooperation between the Mono team and Microsoft. However, since Microsoft wants to see a widespread .NET adoption, it has been quite

supportive towards the project, and helpful in answering many technical queries. If your question was about money... no, there is no cash donated by Microsoft to support this effort.

» .NET is a Microsoft initiative and therefore totally evil because it reinforces MS's monopoly?

The .NET development framework is an open standard with ECMA (www.ecma-international.org) approval. This means that we now exactly what are its various components and how they interoperate.

» Is Microsoft going to release the development framework with an Open Source license?

Maybe when donkeys fly... The Open Specification is conceptually different from the actual implementation. This is one reason why we need Mono: this project provides us Linux users with an Open Source implementation of the .NET development framework.

» Specification... implementation... I'm slightly confused.

Let's give an example: suppose we are following a course on C programming. One of our course assignments consists of writing a program that computes the area of a circle. We both know what we have to do (this is the equivalent of the Open Standard specification), but we have decided to go about it in a different way (this is the implementation part). Nevertheless, the differences are completely transparent to a potential user: were they to use mine or your implementation, users will always obtain as a result the area of the circle. Now, back to the question. The .NET development platform is an Open Standard, Microsoft has developed its own implementation of it (that is and very likely will remain limited to Windows platforms) and Mono provides another implementation, in this case not just Open Source, but also cross-platform (*ie* it can run on different Operating Systems). This is a second reason why we would need Mono: it would bring to Linux (and other OSES) the power of the .NET platform.

» We Linux users would certainly need Mono if we needed .NET; but do we really need .NET?

The answer to this is rather subjective. One fact is that many applications written for Windows in the near future would be written using the .NET framework. In theory, if coherently coded, they should be able to run on any system for which an implementation of the .NET platform is provided. Historically, one of the biggest barriers for widespread Linux adoption has been the lack of powerful apps found somewhere else. We might disagree with this opinion, but for sure it would be an advantage to any newbie if she/he could find exactly

MONO-LOGUE

A very short Mono and .NET jargon survival guide

bytecode intermediate machine-readable code generated by a compiler and executed by an interpreter.
C# (pronounced C-sharp) A clean and powerful object-oriented programming language that is one of the key component of the .NET framework.
CIL (Common Intermediate Language) In the compiler jargon, this term is used to indicate the common languages outputted by the different front-ends of a compiler. Here it is interchangeable with the term 'bytecode'. To be more precise, the .NET CIL is the common bytecode outputted by any .NET-compliant language. Like classical compiler CIL, .NET CIL can be used for optimisations and can be compiled. However,

unlike classical compiler CIL, in our case the compilation process takes place at runtime.
CLI (Common Language Infrastructure) Collection of tools and programs for the execution of .NET bytecode.
CLS (Common Language Specifications) Specifications to which .NET bytecode must conform.
Compiler Program that takes source code as an input and outputs machine-executable code.
Interpreter Program that executes code line-by-line, while reading the source and without compiling it.
JIT (Just In Time) Attribute of compilers that convert bytecode into machine-executable code at runtime.
Runtime engine same as CLI (see above).

the same applications on Linux she/he is used to on Windows. In addition, many developers will find that their favourite development environments will work under Linux, so as well as the applications we will eventually acquire a good number of developers also for Linux-specific projects.

»» Compatibility is a good thing and we all are interested in achieving it. This might well be a good reason for having and using Mono, but I think I'd prefer to carry on using my favourite programming language.

I haven't said you can't use it with .NET, have I? Maybe it is worth to give a look at what the .NET platform aims to achieve. One of the very first choices a developer is faced with is which programming language to use. This should be just matter of taste, but often in this choice the programmer is influenced by the availability of tools, especially libraries.

Do you need to crunch numbers? There are many libraries written in Fortran. Do you want to do system programming? Well, the kernel is written in C... Though is it not impossible to link the executables that are obtained when compiling code written in different high-level languages, it is often tricky and inconvenient to do so. The difficulty in mixing two programming languages resides in the different methods they may use to access variables. This is so because the designers of most programming languages did not care about all the other languages around; maybe they used them as a source of inspiration for the new language, or as an example of procedures the new language should avoid, but rarely did they think about interoperability. This requirement has become a problem only in recent times, when developers may refuse to take part to an Open Source project just because of the imposed programming language. This can be a handicap for larger projects, which generally benefit from attracting as many developers as they can.

In the early days of GNOME and KDE, one of the favourite topics of discussion (and flames!) was whether C++ (chosen by the folks at KDE) was superior to C (official programming language of GNOME), or if it was exactly the other way around. Some developers joined one project or the other exclusively because they had a preference for one language over the other. These two projects sidestep the problem of not putting off developers with strong feelings about the language by offering bindings for many languages (including Perl and Python) to the project official graphical toolkits and development libraries, so we

have *Gtk-python*, *Gtk-perl* and the like. Now, you don't code for a desktop environment because you like the language in which it is written, you can code for it because you like the project itself!

In the .NET framework, this problem just does not exist, since .NET specifically imposes that languages interoperate. If a library has been coherently developed in the .NET framework, it does not matter which programming language its author has chosen, you will be always able to call it from your program, written with your preferred language. »»



WHAT ON EARTH Mono

« What makes such an achievement possible is the inclusion in the Standard of strict procedure intercommunication rules. This by itself wouldn't be a great achievement if it hadn't extended to objects: you can create an object in one language and use it in another language. This is the real good news brought by the .NET platform. What makes it possible is again a new technology called Common Type System (CTS).

»» Which programming languages are .NET-enabled?

In principle, any language. What really matters is the compiler: it *must* output in the format stated by the .NET specifications. Any language with a compiler outputting CIL bytecode can be used with Mono.

»» Ok, I buy it. Mono can be a good thing for us. Why the project is called "Mono" if it is aimed at being "multi"-platform and "multi"-language?

For once, Mono is not an acronym. It is the Spanish word for "monkey". Why monkey? Well, can you think of a Linux company or initiative that is somehow associated with a monkey? Doesn't any image or logo representing a monkey and Linux come to your mind?

»» Surely you don't mean Matt, LXF's increasingly hairy Production Editor? You must mean Ximian!

That's right. Ximian is the company that has founded the project and is backing it up by offering resources and developers. That's the origin of this quite evocative name (don't forget that Miguel de Icaza, one of the founders of the company and leader Mono developer, is Mexican).

»» And why on earth is Ximian sponsoring the project?

Ximian is a Linux company and as such will benefit of any progress of our beloved OS. But Ximian is also a commercial company, and as such can exploit its deep understanding of Mono internals to offer Mono-related services like consultancy and customisation. However, don't forget that Mono is first and foremost a community project: more than 1000 developers are involved at present. This is a huge task force, which goes far beyond Ximian's own resources. Of course, at the moment, the efforts on Mono are mainly coordinated by Ximian people, but the role of independent developers is fundamental in deciding the success and the future directions of the project.

»» How will the recent Novell buyout of Ximian have any influence on Mono's ongoing development?

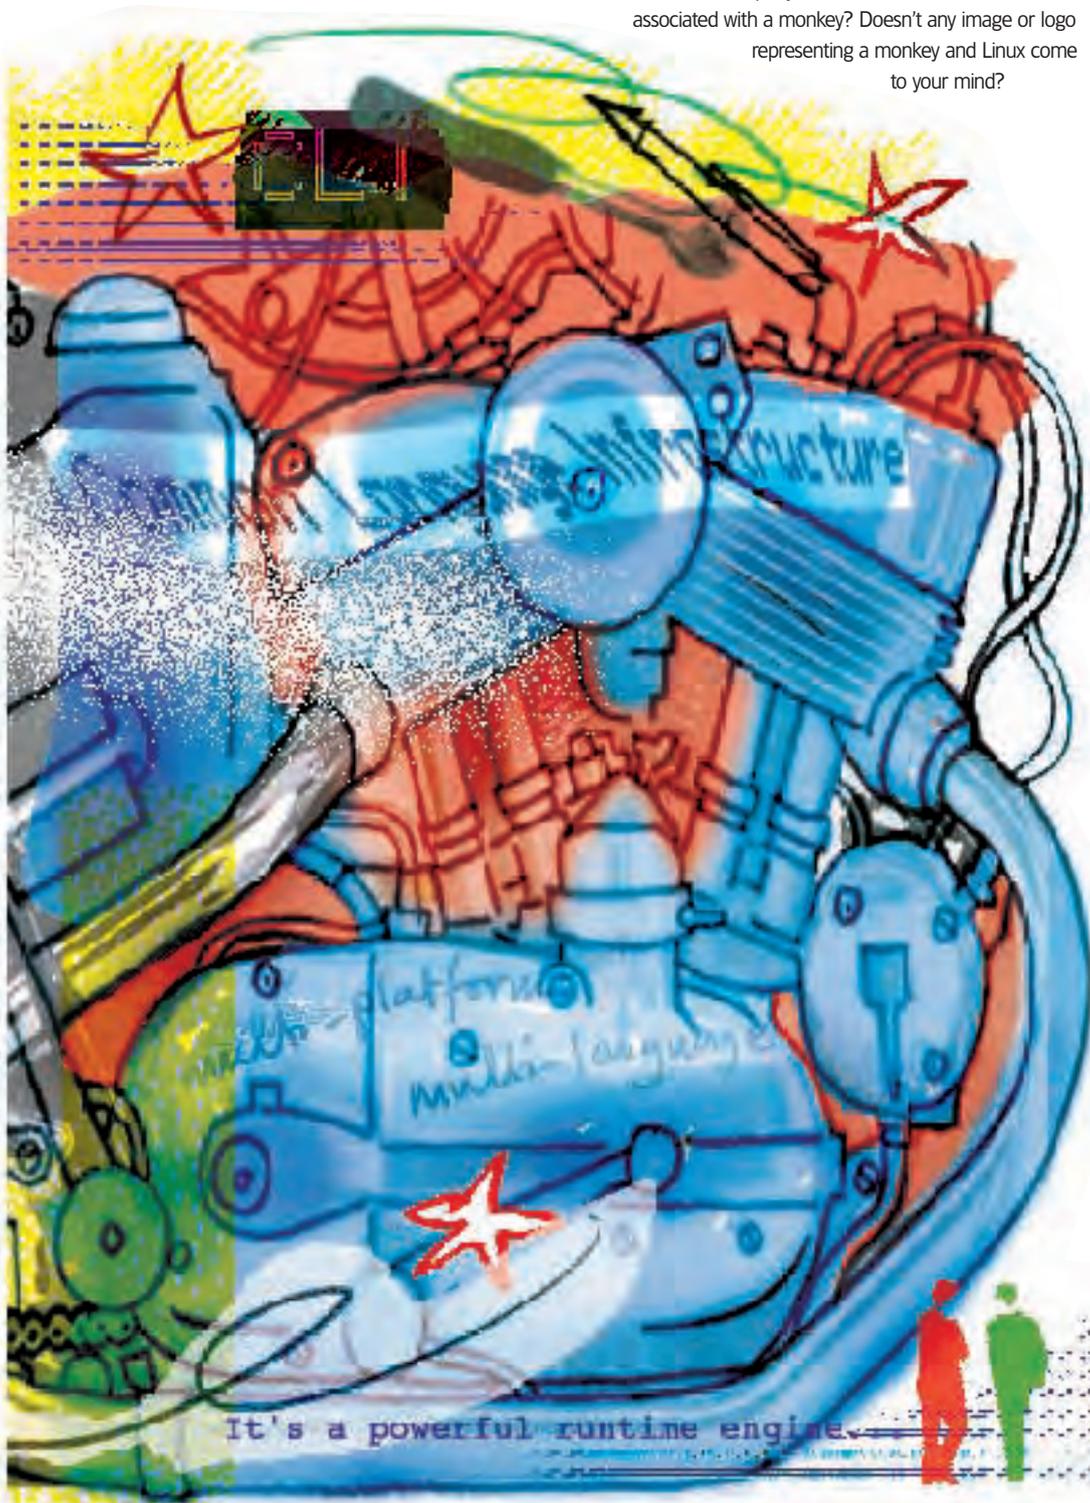
By buying Ximian, Novell has endorsed all the projects and the objectives of the newly acquired company, with the goal of integrating Ximian's technology into its platform. Nothing has changed for the moment; it is easy to foresee a future Novell Linux solution with Mono integrated in it though.

»» .NET and Mono overall sound very promising. I am interested in getting some more technical details. Let's start from the platform itself. What are the components?

The .NET framework is made of three essential components: a .NET enabled compiler, class libraries and a Common Language Infrastructure (CLI) or runtime engine. The chain to build and run an application for .NET consists of the following steps: you start by compiling your application using the compiler for the chosen programming language. The output will then be an 'assembled' file in Common Intermediate Language (CIL). This can be executed by invoking the CLI, which takes care of loading all the required libraries and classes.

»» What are the technologies included in Mono and how does it differ from MS's implementation?

One of the key components of .NET is a new programming language, C#, which is considered by many to be one of the cleanest and most beautiful programming languages around; Mono offers a C# compiler. As well as C#, Mono brings



Visual Basic .NET (VB.NET), ADO.NET and ASP.NET; and many others are scheduled to be ported. This will match what Windows presently offers, which for the moment at least, is more mature.

The libraries imposed by the ECMA standard are already implemented or are going to be implemented, while a few more have been added (often strictly related to the Linux/Unix world, as the POSIX one). Finally we have a runtime engine for the execution of the program.

»» How does the runtime engine work?

The Mono runtime engine implements a compiler for the CIL representation outputted by the high-level compiler (the two must not be confused!), an interpreter for the CIL, threading mechanisms, an I/O interface and a class loader.

»» What is the difference between the compiler and the interpreter?

Actually, I am abusing language somewhat – again! What I am referring to as a compiler is something halfway between a compiler and an interpreter. What it does is take bytecode as an input and convert it in machine code at execution time. The interpreter just misses this last step. The compilation mechanism implemented by the Mono compiler is an example of JIT (Just In Time) compilation.

»» Is there any point in having both a compiler and an interpreter?

The interpreter has the disadvantage of being slower, but works on a broader range of architectures, while the compiler works just on x86 systems. However, the .NET specifications require for the implementation to provide a compiler.

»» What are the supported OSes?

Mono runs on Windows, Linux and Unix. On non-x86 systems, the compiler does not work yet, and the interpreter should be used.

»» Why does Mono support Microsoft OSes, if we can have the original thing on them?

There are several reasons. First of all, Mono is an Open Source project, while Microsoft's .NET implementation is unlikely to ever be anything other than closed-source. Then, you can be absolutely sure that if an application runs with Mono on Linux, it will run with Mono on Windows. Finally, having Mono on Windows can be of great help in establishing how portable the project is in practice.

»» What? So far you have emphasised how portable .NET applications are, yet now you are telling us that there is more than a slim chance that .NET apps will not be properly portable?

.NET applications are fully portable, provided that they are 100 per cent .NET. This means that if only the standard libraries are used, then there won't be any problem. However, different vendors will have different extensions. A program that uses heavily those extensions will not be portable among different implementations. This is why it is important that a single implementation (Mono) supports more OSs, including Windows.

»» How can platform independence be achieved?

.NET programs are not compiled into machine-specific code: they are compiled instead in an intermediate form, which is interpreted or compiled on-the-fly (this is the JIT compilation process mentioned earlier) on the target platform. By not being tied to any machine-specific native language, the program in the intermediate language representation can be executed in any machine with a JIT compiler.

»» This sounds awfully like what Java does. Hasn't Java already achieved all that?

In some way, yes, but the .NET framework plans to go beyond Java bytecode. The language outputted by a .NET compatible compiler must conform the CLS standard. This defines a representation which is a true CIL. The advantage of such a representation is that it allows to optimise the code being compiled.

In fact, the CLS standard mandates some optimisations. Mono adds a few other optimisations to the mandatory ones. Compilers that output bytecode conforming the .NET CLS standard put the target languages halfway between the interpreted and the compiled languages (that's why you often see the two terms interchangeably in this context): like compiled languages, they have a CIL representation for optimisations, but like interpreted languages programs are executed on the spot. The peculiarity is that native code is produced while the program is being executed.

»» How fast is the executable for practical purposes?

Thanks to the optimisation process, the speed of a .NET bytecode executable can be compared with that of a traditionally compiled programming language, with no appreciable performance losses.

»» What about language interoperability?

This is guaranteed by the CLI itself. Once programs or classes have been written in agreement with the CLS specifications, the memory of the underlying programming language is lost and executables and libraries can be easily linked together.

»» So the C# Mono compiler uses gcc?

No, this is an independent compiler, from parsing all the way down to CIL output.

»» To write a compiler from scratch is a big job. Why wasn't a gcc front-end written instead?

Because the two projects have different targets. Moreover, Mono has the constraint to comply to a standard. For this reason, a .NET-compatible C# front-end for gcc would have proved to be more difficult to write than a new compiler from scratch.

»» Why?

Well, the usual compilation chain for a traditional compiler is parsing/lexing>machine-independent optimisation>further optimisation and machine-specific output.

A simple front-end for gcc wouldn't have worked because at no stage in this chain does gcc output .NET CIL. So both a back-end and a front-end for gcc should have been written. Moreover, the ECMA standard also specifies which optimisations have to be performed, and this would have reflected also on a modification of the intermediate step.

Another difficulty is that the language chosen for the C# compiler was C# itself, so this would have created the usual chicken-and-egg problem (to build a C# compiler for gcc you need a C# compiler for gcc). Nevertheless, it is possible to build a C# compiler integrated in the gcc collection.

»» And writing a gcc back-end targeting the .NET CIL instead of the machine-specific object file?

This is a more interesting option. Of course it is possible and it would be desirable, since this will give .NET bytecode for all the languages supported by gcc in one go. However, nothing has been achieved yet in this area.

»» Is there any relationship between GNOME and Mono?

Mono project leaders are former outstanding GNOME developers. It is natural then that GNOME is one of the targets of the Mono offer. In fact, one of the .NET Mono-specific projects is Gtk#, which provides C# bindings for *Gtk*, *gnomelibs* and other basic libraries of this desktop environment.

»» How usable are the Gtk#?

At such a point that you can write fully functional GNOME apps with them.

»» Will GNOME be based on Mono?

It's still early days for the Mono project. Only time will tell. However, Mono has the potential to become a standard development environment for many Open Source projects.

»» What about KDE?

Qt#, which has its project-page at <http://qtsharp.sourceforge.net> is an effort by the KDE community to provide C# bindings for the Qt libraries.

WHAT ON EARTH Mono

Are there any IDEs to be used specifically with/for Mono?

Mono tends to privilege C#, so what you are looking for is actually a C# IDE. Currently, Mono developers can choose among three IDEs:

a Good old Emacs with a C# Elisp Mode (get it from www.cybercom.net/~zbrad/DotNet/Emacs or <http://dabh.dk/script>). Emacs has the advantage of being installed or installable on any Linux system; however, it is not specifically an IDE and so you might also have to rely on external tools.

b Eclipse with the C# plugin (www.improve-technologies.com/alpha/esharp). Eclipse is a powerful multi-purpose and extensible IDE; on the minus side, it depends on Java and installation may be not straightforward.

c SharpDevelop. This is a .NET IDE for the .NET framework. Unfortunately, it is not much of an option for Linux at present, because it builds on the class System.Windows.Forms, not fully functional at present on Mono.

Any other development tools?

Mono has a powerful graphical debugger.

Can I write GUI apps on Mono?

There are several possibilities, but probably the easiest one is to use Gtk#. Though the GUI class specified in the standard, System.Windows.Forms should be in principle preferred for a 100 per cent .NET app, at the moment its implementation is not complete.

Can I build Web apps on Mono?

Yes. Mono supports ASPNET (but not v2.0 yet). What characterises this language and makes it different from all the other Web-oriented languages is

a wide collection of objects that developers can painlessly embed into their pages with just one instruction. In this sense, ASPNET is object-oriented. Like object-oriented languages, ASPNET can be extended by developers with new objects. ASPNET can be used for web pages as well as for web services.

Does Mono offer database interconnectivity?

Yes. ADO.NET (shipping with Mono) is a class library aimed at exchanging data with relational databases.

What is the status of the project?

Though it is still considered beta software, it is quite usable, and it has been this way for a while. The software is mature to such an extent that Mono has been self-hosting for about a year and half by now. Also, the classes and libraries are almost complete. The C# compiler is essentially ready, and the runtime engine works pretty well. There are a few issues with some supported programming languages and APIs, with Basic.NET and ADO.NET still lacking a bit behind, and some issues of compatibility with Windows-tailored libraries like System.Windows.Forms, but none of these limitations should be a major problem.

When will Mono 1.0 be released?

The new roadmap for the Mono project has been just posted. It is targeted for release in the first quarter of 2004. This release would conform the 1.1 .NET API specifications while offering compatibility with the 1.0 specifications. However, the general understanding is that only the components that would be considered ready by the deadline would make it into 1.0. All the other components (including the Visual Basic compiler) would be put into a mono-unstable package.

What about the future?

Mono has been self-hosting for about one year and half by now. The project will be kept aligned with the .NET API specifications and will bring new mature technologies with the stable releases. For instance, the Visual Basic compiler should be part of the stable package in version 1.2, which would conform the 1.2 .NET API specifications.

Will I be able to write proprietary applications under Mono?

Yes, Mono license allows you to do that.

Which license does Mono use?

Different components are covered by different licenses. For instance, the C# compiler is released under the terms of the GNU GPL, the runtime engine is covered by the GNU Library GPL and for the class libraries the terms of the MIT X11 License have been chosen. Mono can be licensed under different terms and conditions by Ximian to those who can't use the GPL License.

I want Mono now. Where can I get it from?

From the download section of the Web site of the project, www.go-mono.com. Here you will find binary packages for a few distributions and the source code package. Even if your distribution is not listed here, the chances are that you will still be able to find in your installation CDs a prepackaged version of Mono. If you plan to install from the source code, it would probably be a good idea to read the detailed installation instructions given in the Mono Beginning Howto www.go-mono.org/mono-beginning.html.

As is usual for Open Source projects, it is also possible to have the latest unstable release either by installing from a daily snapshot package or by checking out the sources from CVS. Again, read the documentation carefully if you plan to go this road.

And then how do I use it?

It depends on what you want to do. Suppose, for instance, you want to write and run a C# program. First, with your favourite editor or a C# IDE you write the code and save the corresponding file (called for instance "hello.cs"). Then, you compile it with

```
# mcs hello.cs
```

and finally you run it with

```
# mono hello.exe
```

Where can I find more Mono info?

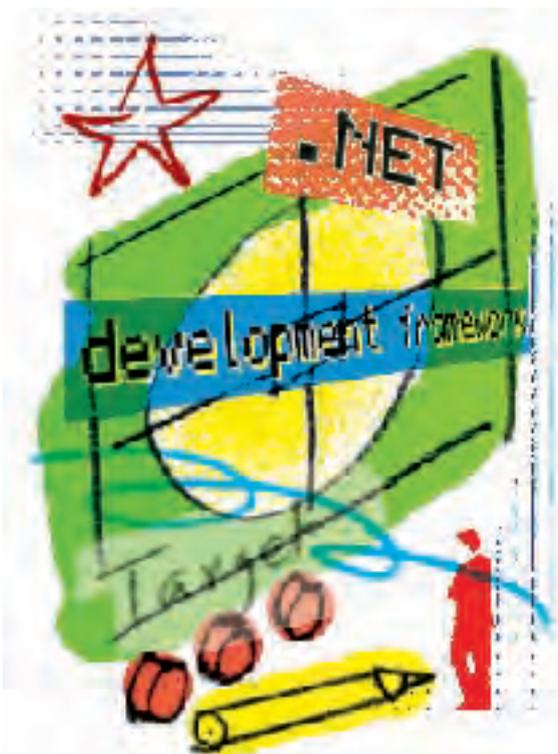
The web site of the project is a good place to start. It is very well detailed and contains in-depth descriptions of both Mono (for users) and its internals (for developers). It also contains useful links to .NET-specific information, including the ECMA standards. Another good stop is the Mono community site www.gotmono.com. Among other things, here you can find a lively forum and the *Mono Handbook*. More than an organic document, the latter is a collection of contributions from volunteers covering many aspects of the project. It is partially incomplete, but contains good descriptions of several aspects of Mono and .NET.

Is Mono the only Open Source implementation of .NET?

No. There is another Open Source project called DotGNU (<http://dotgnu.org>). These two projects have very similar objectives.

Are the two projects cooperating?

Sadly, no; however, this is not as bad as it might seem. One reason for it is that the Open Source license guarantees reusability of the code, so there is not really duplication of effort; another reason is that competition stimulates innovation. Thanks to competition between KDE and GNOME, today we have two first-class desktops. Thanks to the competition between Mono and DotGNU, we might have one day the best implementation around of the .NET development platform. 



Tutorials

Our experts offer help and opinions on a whole host of Linux applications

YOUR GUIDE TO GETTING THINGS DONE!

Whether you are just starting out in Linux, or an experienced veteran, there's always more to learn. Every issue of *Linux Format* is packed full of practical advice, and nowhere is it more concentrated than in our tutorials pages.

Here you'll find expert guides to all sorts of things, from Basic Linux usage to understanding and deploying network solutions, from simple script coding to the complexities of Perl regular expressions, Java server apps and more. We aim to bring a good mix of tutorials to each issue, but if you have any suggestions for topics you'd like us to cover, why not contact us, by email at linuxformat@futurenet.co.uk or by snail mail, or log on to www.linuxformat.co.uk and post your suggestions in our special forums? Hope to hear from you soon!

Nick Veitch EDITOR

HOW CODE IS REPRESENTED

Including code in magazines can be tricky, but we hope our notation will help it become clear. When lines are too long for our columns, the remaining text appears on the next line in a solid blue box:

procedure

```
TfrmTextEditor.mniWordWrapClick
(Sender: TObject);
```

otherwise, there is usually a gap between lines:

```
begin
```

```
  mniWordWrap.Checked := false
```

```
end;
```

Usually, you'll find the code on our CD/DVD too.

THIS MONTH TEACH YOURSELF...

Beginners' Guide: Ximian Evolution

Sort your entire life out with the PIM that's the GNOME equivalent of KDE's *KOrganizer* **p62**

GIMP 2 preview

Changes and new features in the latest iteration of Linux's definitive art app explained and rated **p66**

Kino >>

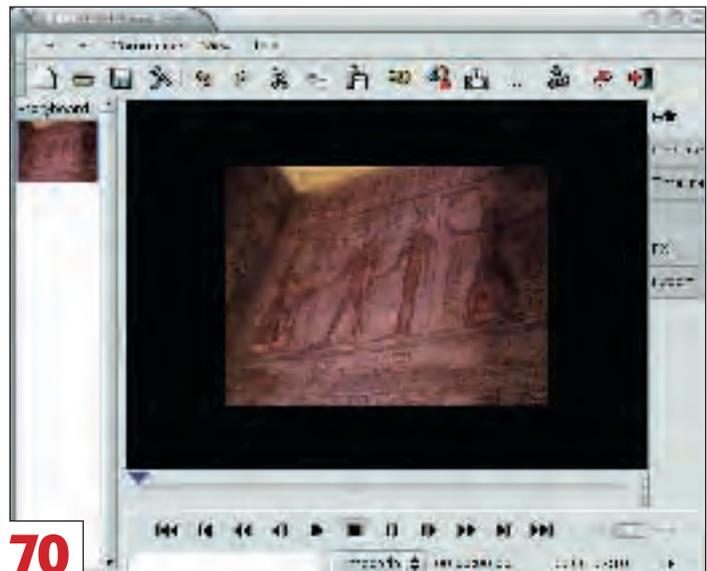
A quick guide to joining your video clips together using professional-looking effects **p70**

Practical PHP

PHP has nominal SNMPv3 support: so what are the seven available functions, and what can you do with them? **p72**

Server School

Database servers – don't confuse the differences between *MySQL* and *PostgreSQL* **p76**



70

Perl Template Toolkit

Generating configuration files needn't be a nightmare – soothe any hiccups in your processing runs with a good dose of *ttree* **p80**

SNMP

Here come the MIBs – how the Open Source NET-SNMP tools make your network management truly a simple affair **p84**

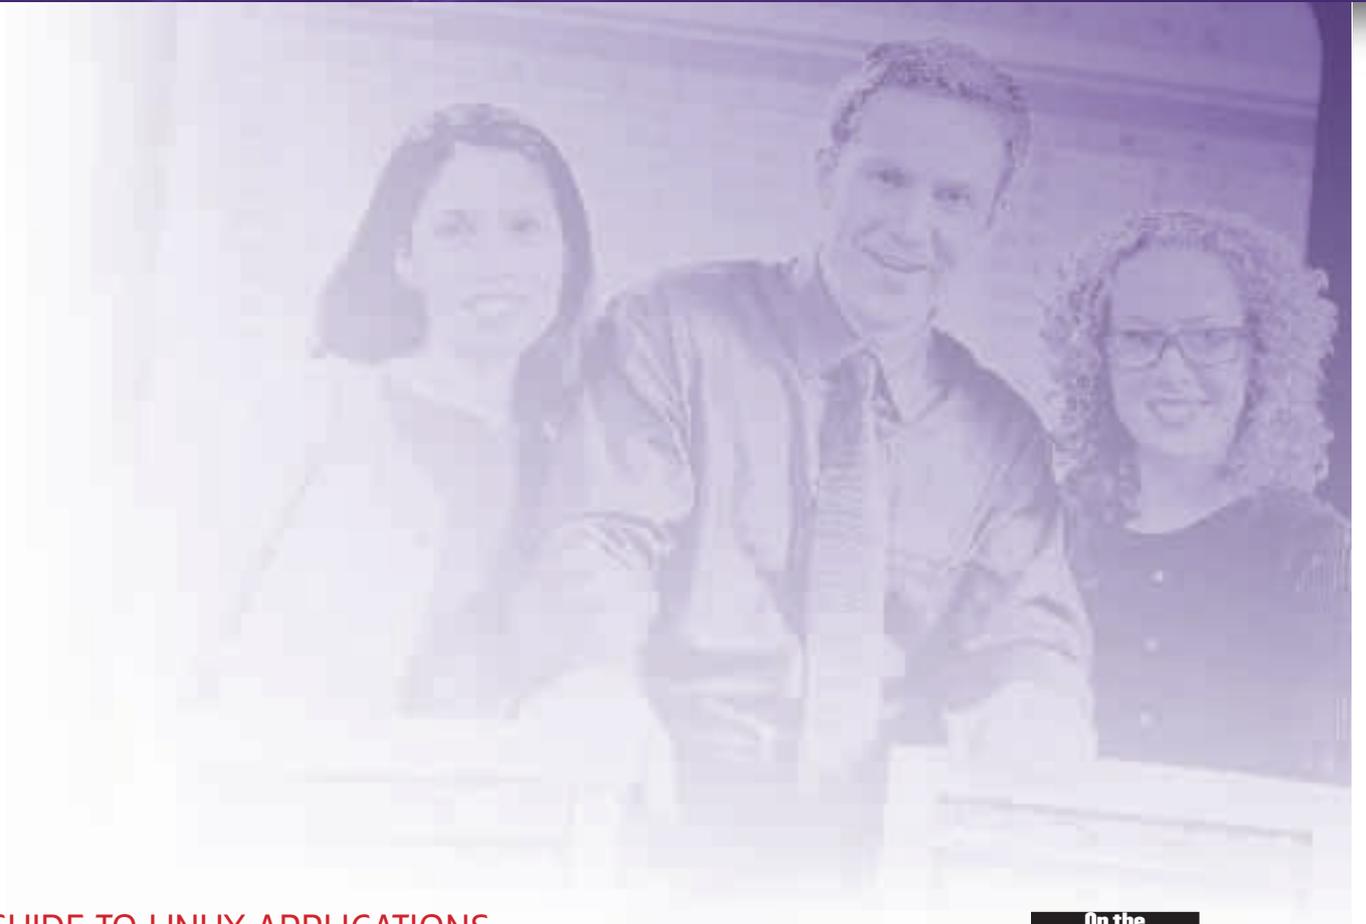
TIP OF THE MONTH! KEEP IT UNREAL!

On this month's DVD coverdisc, we've supplied the *Unreal Tournament 2004* demo, which works almost as well as the full version reviewed on page 22. As the code for the demo was released a few weeks before the final build, there are a few more bugs in there than in the release code, but that shouldn't stop you getting into a frag-fest with anyone else you can tempt into your sights.

The one drawback to installation is that you need a fairly high-spec computer. At the absolute least – and this is if you turn all the detail down and even then don't mind it

being a fairly poor gaming experience – you need a 600MHz CPU with a first-generation GeForce. Anything beyond that, particularly in the graphics card area, should be OK. The official minimum requirements call for a 1GHz CPU with a 64MB GeForce 2 or higher – your mileage may vary. The biggest known issue in this version – and even in the final release – is fairly slow sound processing, so if you don't have speakers plugged into your system for some reason, add **-nosound** to the command-line to disable audio and get a speed boost. For more information on the full

game, as well as how to buy it once you've finished fragging your way through the handful of levels in the demo, check out our full review. There are also only six game types included (Assault, Onslaught, Deathmatch, Capture the Flag, Team Deathmatch, and Bombing Run) compared to the ten in the full release. Note that the demo version on our discs is v3120 – newer versions may well be available online by the time you read this, and, if you can handle the 200MB download, a newer version is well worth getting: <http://www.unrealtournament.com/ut2004/downloads.php>



BEGINNERS' GUIDE TO LINUX APPLICATIONS

Sort your life out with: Ximian Evolution



To balance last month's Kapplication overload, Andy Channelle introduces the GNOME alternative to *KOrganizer* – Ximian's excellent *Evolution 1.4*.

Last issue we dissected *Contact/KOrganizer*, an application for KDE that integrates calendar, news and email functions in one handy package. But what if you've opted for the splendid GNOME desktop environment or have an allergy to all those apps beginning with K? Well, the answer could be *Evolution*, the Outlook-style personal information manager

(PIM) created by Ximian, which, again, pulls a number of jobs such as address book, email, calendar, news and weather services into one consistent user interface.

The current version of *Evolution* is 1.4, though the 1.5 testing phase is underway (which will lead to the eventual release of *Evolution 2*), so we will also try to highlight some of the features expected in the next release. Chances are that *Evolution* is either installed as standard or will be accessible from your distribution's CDs (or this month's *LXF* coverdiscs). If not, Ximian has developed a network installation option which should work with most recently released distributions and, obviously, depends on your having a active (and preferably broadband) Internet connection. To get *Evolution*, open a terminal and use **su** and the computer's root password to gain administrator privileges. Next type

```
wget -q -O - http://go.ximian.com |sh
```

The pipe symbol, "**|**" is accessed by pressing **SHIFT+|**. This launches Ximian's installation application and offers a choice of installing *Evolution* on its own, or as part of the full Ximian XD2 desktop (which also includes a themed version of *OpenOffice.org* and other software). Simply follow the prompts to "download from an official Ximian mirror" (unless you have a Red Carpet account) and, if necessary, registering yourself as a Ximian user.

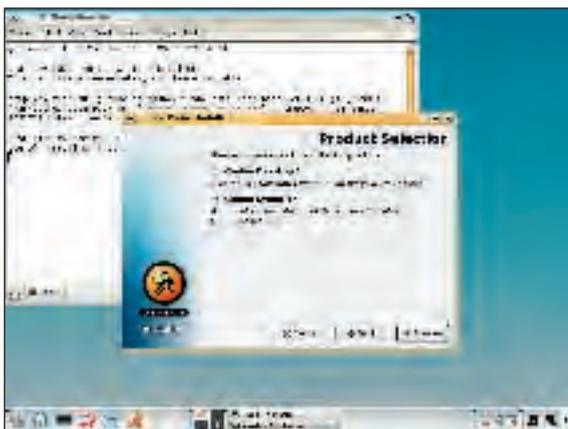


Fig1 Ximian's network installer requires root access.

Occasionally we had difficulty running the installation application from within a user account, so if you have any problems, just log out and then log in as root to complete the job.

Once *Evolution* is installed, we should be able to access it from the GNOME/KDE panel, or by selecting 'Run application' from the menu (or opening a terminal/console) and typing **evolution**.

Upon launch, and following the mail wizard which provides options for setting up a default mail account, we are presented with the 'Summary page'. As we discover more of *Evolution's* features, this page will assume an important role in our day-to-day computer use – which is a shame as the next generation *Evolution* release does away with it – to provide an overview of events and tasks that are on the horizon, emails waiting for attention, the latest news and even the weather. At first sight, it may all seem a little generic, but we can change that (see *The Daily Me*, on the following pages).

The Day Today

In this tutorial we're most concerned with *Evolution's* skills as a calendar and task management application, and like any information system, getting the best out of it depends on putting clear and concise data in. In short, the more comprehensive our input, the better we can filter output to our needs. This is important because, until the next release at least, *Evolution* cannot 'overlay' various calendars as we did with *Kontact* last time, so careful categorisation is the watchword this month!

A. The interface has the by now standard combination of menu bar and button bar with various viewing panes beneath them. At the far left of the button bar is the 'New' button. Not only does this provide tools to add an appointment to the calendar, as expected, but we can also add contacts, send mail and add new folders to *Evolution's* hierarchy. We will start by adding an ordinary appointment, so hit the 'New' button; we can click the down arrow next to this button to get other options, but we are just creating an appointment at this time. The standard Appointments dialog should appear ready for our data (**Fig3**). In this dialog, the Summary is what will finally appear in the calendar and we can also add a location, start and end times of the event (or make it a non-specific/all day appointment), add notes to the entry, set a classification (really only necessary if you're using *Evolution* as a full groupware client) and define a category, as is illustrated in **Fig4**.

As mentioned earlier, the categories are vital when using this application. We click on Categories to open the dialog and either choose one of the available ones, or hit 'Edit Master Category List.', to add, edit or delete entries. Once that's all done, select the reminders tab and set *Evolution* to either show a message, play a sound or run a program to alert you when the appointment is imminent. To display a message 15 minutes before the meeting, we select the relevant parts from the drop down lists, hit the options button to input the message we want to display and then hit 'Add'. Usefully, we can also set a message to be displayed at a defined time after the appointment, so if, for instance, you wanted a reminder to email participants a list of goals or outcomes from a meeting, you could set a reminder for this purpose.

If you are organising a meeting, it is a good idea to set this up within *Evolution*, so that changes to venue, time or participants can be quickly communicated to the rest of the group. Simply hit the 'Meeting' button in the 'Appointment' dialog to add names, email addresses and functions to the meeting. Once this is done, every change made to the entry will result in

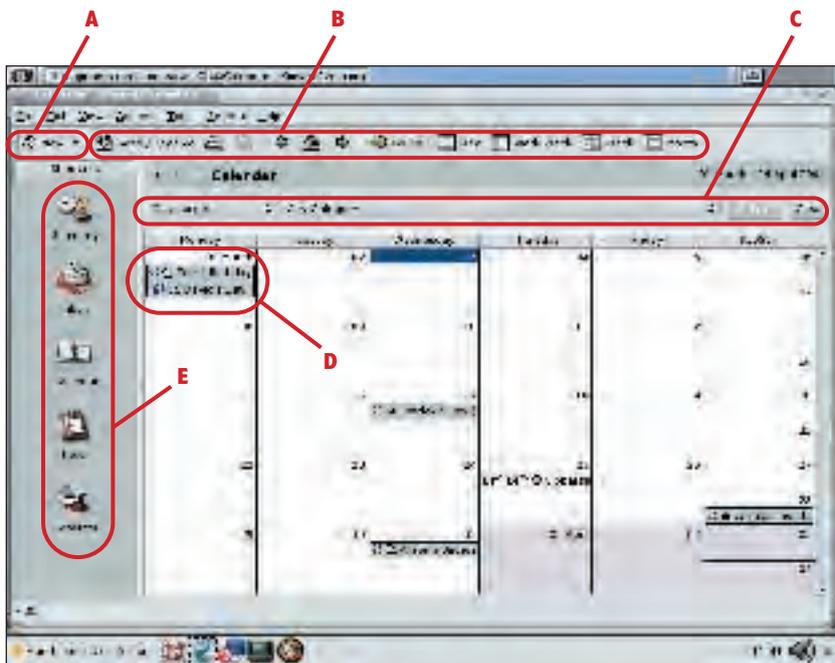


Fig2 **EVOLUTION 1.4:** The interface explained – see the lettered explanations in the text.

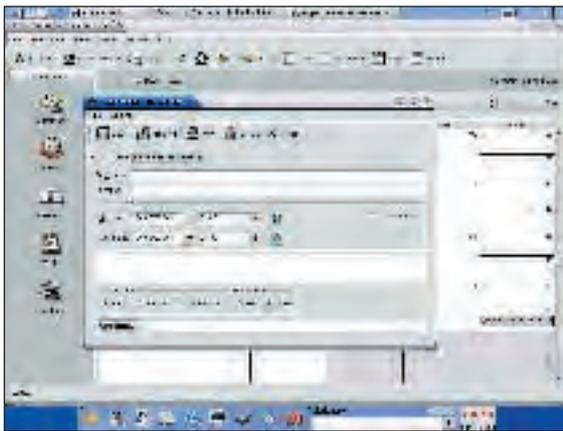


Fig3 Appointments can be entered as one-offs or recurring.

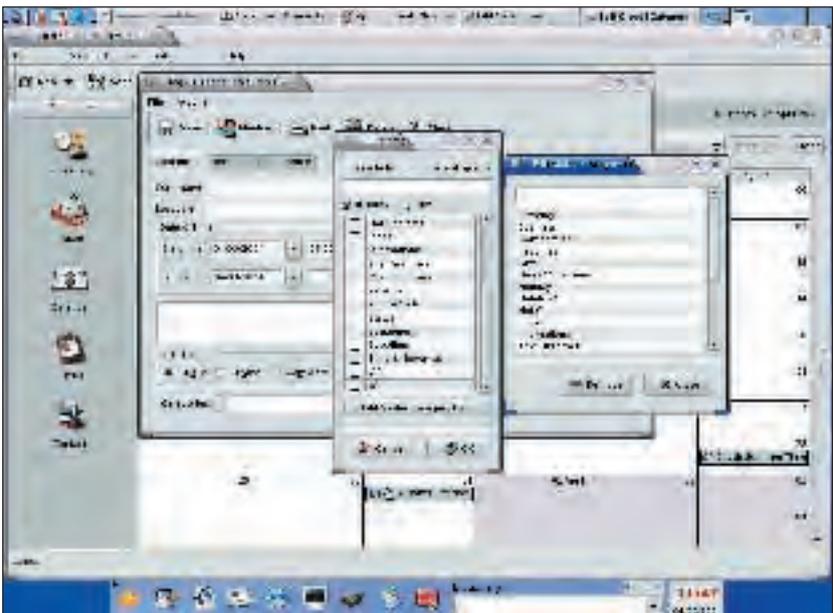


Fig4 Choosing categories may look a bit daunting to start with, but gets easier with use.

TUTORIAL Beginners' Linux: Evolution

◀◀ *Evolution* offering the option to send an amended schedule to all participants. Actually amending an entry is as simple as double clicking on it in the calendar to launch the 'event' dialog, making the alterations and hitting 'Save'.

The 'Recurrence' tab allows you to add periodic dates (anniversaries, birthdays, etc.) to the calendar. Simply select the frequency of the entry and its end date, if appropriate.

B. The rest of the toolbar consists of icons for sending and receiving queued mail, printing, deleting entries and a set of navigation tools. From left to right, these are Go Back, Go To Today, Go Forward and Go To a specific date. This last will open a calendar widget for visual navigation. Finally there are a series of views allowing us to view a month, week, working week or day at a time. If you're a part-timer or work shifts, you can reset the 'work week' to suit your hours. Do Tools>Settings and choose the 'Calendar and Tasks...' option. Here you can select which days

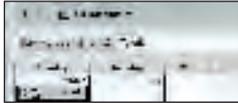


Fig5 Widgets help filter your information by one or more categories.

ICAL AND EVOLUTION

Importing calendar information

Just like *Kontact/Korganizer*, *Evolution* is capable of importing data in iCalendar format (with the .ics extension) which means there are thousands of calendars out there on the Internet covering almost every conceivable need. All we need to do is find it, import it and enjoy. To do this, go to File>Import, select 'Import a single file...' and browse to the file you want to import. Be aware, though, that information imported in this way is merged with the current calendar – and there's no easy way to disentangle it – which is why we stressed the need for clear categories earlier.

While *Evolution* can't yet manage to overlay calendars in the manner of *Kontact* (it is a feature due in the next release), we can add a new folder to the application which can then be viewed as a separate calendar. So if, for example, you are obsessed with the lives of the rich and famous, you can download an almost

daily updated iCal file of celebrity birthdays (from <http://ical.mac.com/tsmallwood/Celebrity.ics>) to ensure you never forget to send Roger Daltrey a birthday present again.

Do File>New>Folder...; give it a name, designate it as a Calendar file and set it as a child to your main calendar. Hit 'OK'. To access the new folder, ensure the 'Folder View' is active and then simply navigate to the calendar in the normal way. You should now have a fresh calendar waiting for your instruction, and you can import the new file here.

You can see the 'default' calendar by clicking on its name in the folder view, and deleting a calendar is just a case of right-clicking on its name to bring up the context sensitive menu and selecting delete. Admittedly, it's not as elegant as *Kontact's* method, but it is workable, and some people may find the clearly defined distinction between calendars a blessing.

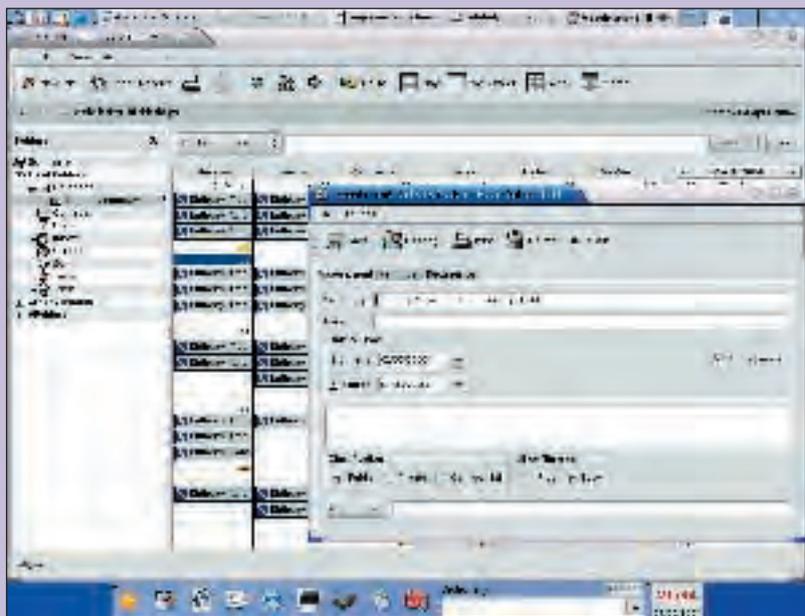


Fig6 Take care not to import any iCal calendars over the top of your existing ones...

should be displayed as Work Days, and also define a start and end point of your hours. These hours are then highlighted in white, as opposed to grey for the rest of the day.

In the day view, events which are scheduled for all day, usually birthdays and so on, will appear at the top of the 'page' while timed appointments will appear at the appropriate time – there's no *Kontact*-style timeline ticking away on the UI though, which some people may prefer.

C. These are the all important filters. Busy people, and those who just like acres of pointless information at their fingertips, can soon build calendars that are visually overwhelming, and the information vital to one day's task may get lost among the 'noise'. While a sensible alarm system can do much to alleviate this, it is much better to organise things so you can, say, view only work related entries, or just see birthdays. However, this section is very versatile and goes beyond merely removing from site entries based on a category.

The first widget allows us to select whether to filter on entire calendar entries, the summary, description, comments or the category itself, while the second widget provides a space to define the search string. So, to find all 'Tutorial' deadlines, we would select 'Summary contains...', then add the word Tutorial in the search space and hit the 'Find now' button. Everything that doesn't contain this string is temporarily removed from the main view. To get back to the complete picture, as shown in Fig6, select the 'Clear' button,

By the way, just above this section, you can click on the word 'Calendar' to access all the folders associated with *Evolution* to provide quick access to things like sent mails, trash and contacts (it's also possible to opt to have this displayed permanently with View>Folder Bar) or cycle through the available applications.

D. These are the basic calendar entries, and they have icons to denote their category – if assigned – and whether they are recurring events or not. While not a feature of this release, *Evolution 2* will have a user definable colour scheme allowing us to set birthday's in red, work deadlines in blue etc, which should both improve useability and make things a little less austere. At the far right of this pane, there is a moveable bar; clicking and dragging the widget (denoted by three small dots) will open up an extra pane with a clickable 'small' calendar view and a list of upcoming tasks.

E. This is the main selection panel, where we can choose which part – mail, contacts, calendar, etc – that we want to see in the main window. If the order of these entries doesn't suit the way you work, you can adjust it by simply dragging and dropping icons to a difference location. You can also rename or even remove entries from here by right clicking to pop up a context sensitive menu. If you would rather not see this shortcut bar, you can remove it by doing View>Shortcut Bar.

Task inspector

While email and contacts are the glamour girls of the communications revolution, the humble task manager – derided as a ToDo list by some – is often left unused. This is a mistake, because beyond simply showing us when something is due, a decent task manager enables us to prioritise projects based on deadline, importance and completion. However, the default task screen in *Evolution* is very disappointing (Fig7), and to make the most of it, you will need to spend a short time setting things up.

As standard, we are shown just the name of the task and a checkbox to show whether it's completed; that's not good

enough! To remedy the situation, do View>Current View>Define Views... to launch a small dialog. Now do 'New' to start a new view, give it a name and then, after highlighting 'table' in the second section and hitting the OK button, select 'Fields Shown...'. It is now possible to add a range of options to the task screen, including useful things like the due date, percent complete, etc. These are ranged down the left of the screen and we add them to the view by selecting and hitting the 'Add' button to add the field to the list on the right. You can add and remove and move fields up or down (which equate to left-right on the final screen) until you find the right balance. Finally, 'OK' through the remaining dialogs and, once back on the main screen, do View >Current View> and select the view by name (**Fig8**).

The Daily Me

Nicholas Negroponte, – founder of Wired magazine – once highlighted the possibilities of personalised newspapers delivered to your inbox every morning; a good idea, but *Evolution* goes one better and provides facilities for a constantly updated, bespoke news service using the magic of Rich Site Summary (RSS).

To configure news sources, first go into Tools>Settings and make sure the 'Summary Preferences' section is selected. Along the top of the main window are a series of tabs, and we need to choose News Feeds. This section is divided up primarily into two panes: on the left are potential sources of news, and on the right those we have selected. To add a new source, simply highlight it on the left and click the 'Add' button. Similarly to remove one, highlight it on the right and click 'Remove'.

Just below the lists there is a pair of options to define how often source sites are polled for new items, the Refresh time, and how many items should be shown. Busy sites such as Slashdot don't appreciate being contacted every five seconds for an update (it consumes bandwidth and is quite pointless anyway) so it is best to stick with a minimum of 10 minutes, or 600 seconds, between refreshes.

The range of news sources pre-defined in *Evolution* is quite limited, mostly covering computer/Linux subjects, but you can add any RSS feed you want. For instance we can add a feed from the BBC News front page. Hit the 'New Feed' button and add the following details. Name: BBC News. URL: http://news.bbc.co.uk/rss/newsonline_uk_edition/front_page/rss091.xml. Now hit **Enter**, and then add the new feed to the page. You may need to 'Reload' the page before the changes take effect (**Fig9**).

RSS is quite a hot topic at the moment and many news sites are beginning to offer the service, though there is still a little confusion over the name. If you are browsing and find a source of interest, look on the page for a small button marked with either 'rss' or 'xml'. Clicking on it – or right clicking the link and selecting something along the lines of 'copy link' should allow you to then paste it directly into *Evolution*.

We can also add one or more weather services to the summary page to get a reading of the current temperature, wind-speed and atmospheric pressure from locations around the world. Visit the Tools>Settings>Summary Preferences and this time select 'Weather'. The configuration tool is consistent with the news feed one, so just select the location you want to see in the left pane, click the 'Add' button and select a refresh time. The list of available locations is very comprehensive, and we are not limited to one. Once the weather is set up, clicking on My Weather in the summary should open the web page of your local weather service; in the UK it links to www.metoffice.gov.uk. 

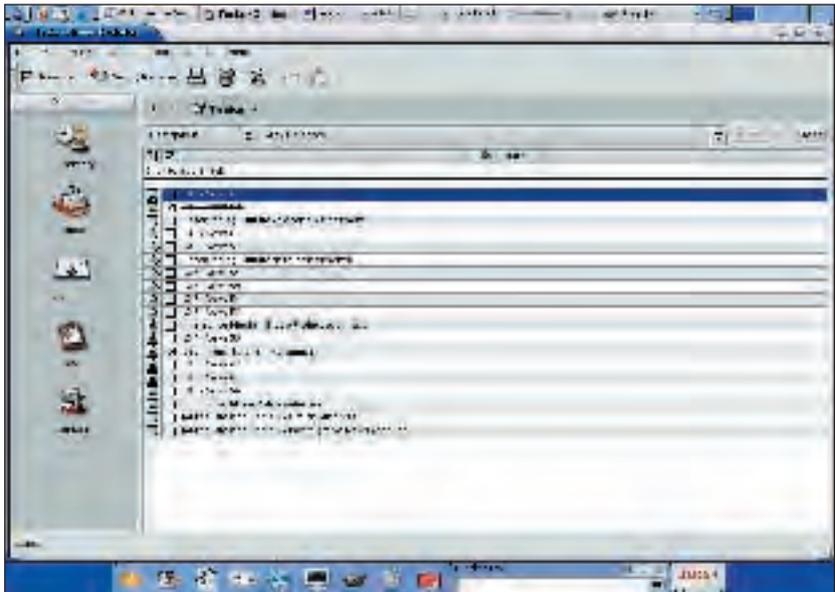


Fig7 The Task Inspector summary is often overlooked, but is a boon for busy users.

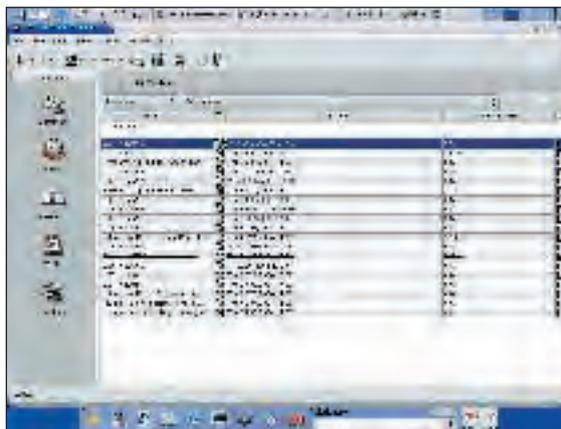


Fig8 The Task Inspector summary with all completed work struck through – time to go to the pub!

NEXT MONTH

The fantastic Open Source audio editor *Audacity* has just been updated, so we'll be checking out some of the new features of version 1.2.

If there's an application or issue you think we should be covering, please visit the LXF forums (www.linuxformat.co.uk) or send an email to andy@channelle.co.uk to let us know.

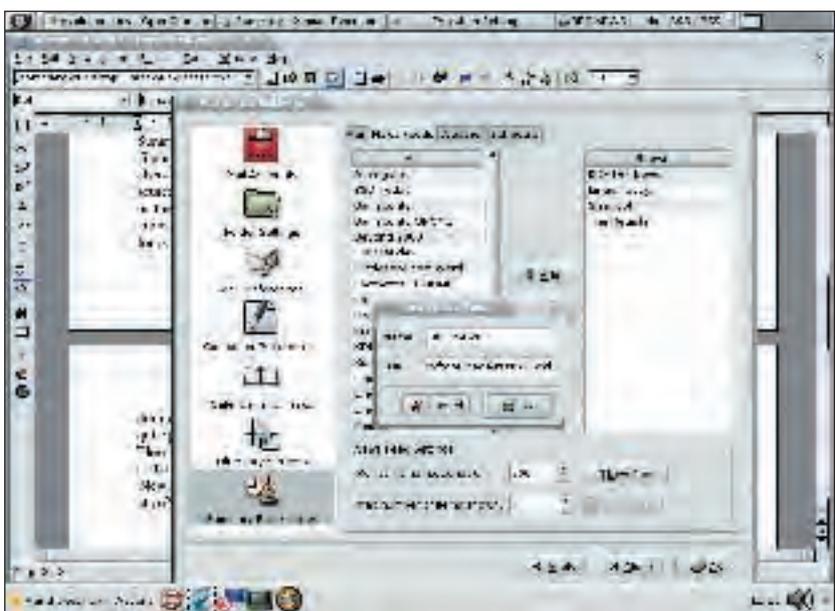


Fig9 RSS delivers up-to-the-minute news direct to your desktop.

TUTORIAL GIMP



NEW VERSION UPDATE

The GIMP 2.0 Preview



After more than three years of development, the next major step in the evolution of Linux desktop graphics is about to be revealed – and it has been worth the wait, as **Michael J Hammel** reveals...

The GNU Image Manipulation Program – better known as *The GIMP* – is about to take another major step forward. The 2.0 release, which has been in development since late December 2000, is being readied for mass consumption by the small band of volunteers which keeps the red, green and blue blood flowing in this popular project.

The GIMP 2.0 comes with a huge list of improvements, from cross-platform support to a user-friendly design to enhanced text and vector drawing support, and much more. In this article we'll look at what's new, what's improved, and whether they got it right after all this time. To cover it all, we need to start at the most obvious change: the user interface.

DOCKS AND UI

The new UI

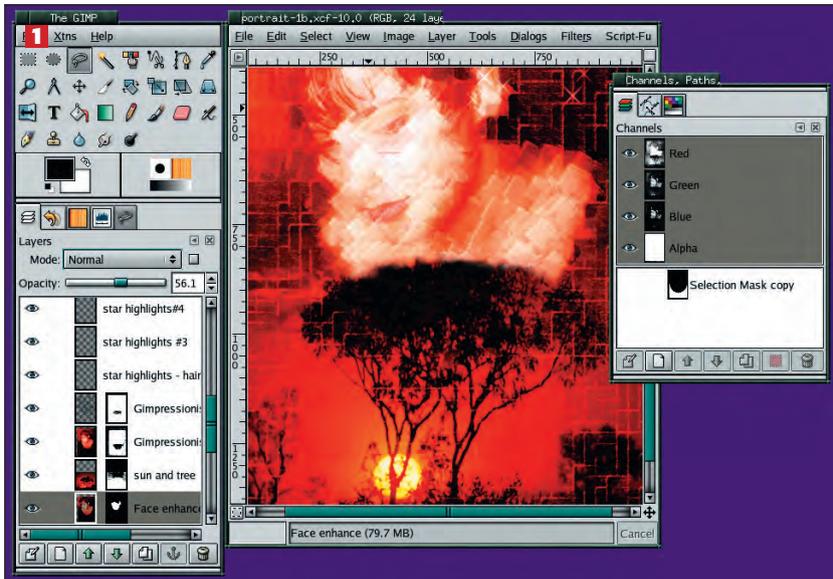
1 Users of the current release of *The GIMP* have long complained about two big issues with the user interface. First, it just wasn't obvious how to use the *The GIMP*. Too many features were hard to find. Most new users stumbled blindly before ever finding the Canvas menu (right mouse click in an image window).

Second, *The GIMP* does more to clutter a desktop than any open source tool around. Too many windows. So little space. Managing your work environment was never an easy thing with *The GIMP*.

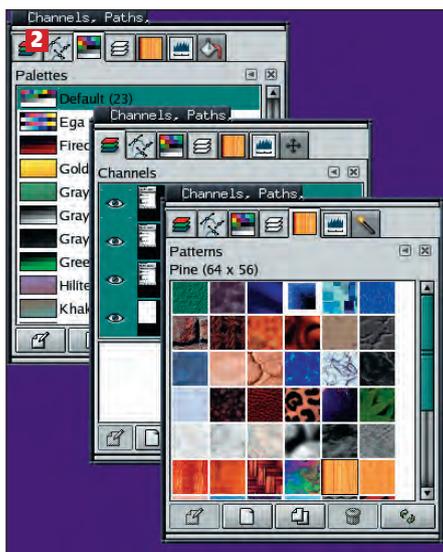
Both of these issues had modest hacks provided during the lifespan of the 1.2 release. This time around, the *GIMP* Developers took a much bigger stab at it.

The 2.0 user interface gets much of its improvements from the *GTK+ 2.x* toolkit it uses. *GTK*, the *GIMP Toolkit*, was born with *The GIMP* but has long since become a powerful entity in its own right. This toolkit allows *GIMP* to 'dock' its windows – grouping them in both stacked and tabbed formats – to address the desktop clutter issue.

The GIMP's new user interface is both polished and easy to use.



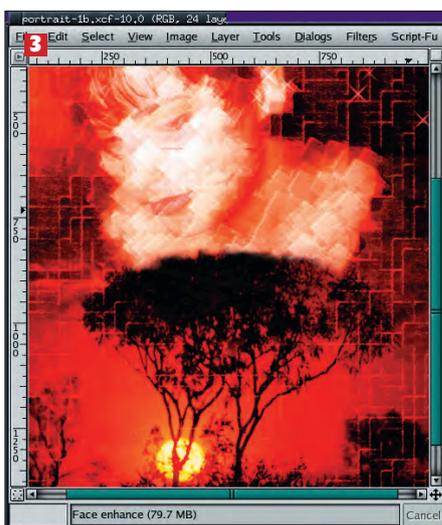
GIMP Docks allow multiple windows to be opened in a single, tabbed interface.



GIMP Docks

2 Docking provides users a method to manage their workspace more effectively, permitting dialogs to remain open and quickly accessed without taking up additional screen space. And there is no limit to the number of docks or stacks allowed. Add to this the new full-screen mode, which can be configured with menus, scrollbars, rulers and other options independent of the normal view mode. The wealth of configuration here is nearly endless.

This new-found wealth of options also leads to the solution for the other big UI problem – ease of use. Menus now adorn the top of Canvas windows by default. These menus, like most features of the user interface, can be hidden from view, opened with the Canvas Menu button (the right arrow in the upper left corner of the Canvas window) or even accessed using the old right mouse click, if desired. Menus are better organised now as well, with layer features in their own menu, colour in their own, and so forth.



GIMP Canvas

3 Going beyond menus, the Tool Options dialog has improved as well. The Selection tools, for example, all provide buttons to change modes between adding, replacing and subtracting. This feature replaces the overloaded use of **Ctrl/Shift/Alt** key and mouse-click combinations required in *The GIMP 1.2*. And the paint tool options now all come with their own brush and pattern settings. The paintbrush and pencil can use different brushes without you having to change the brush after switching tools. Better yet, the brush, pattern and gradient options can be selected using a mouse wheel in the Tool Options window!

The new GIMP Canvas window makes menus obvious and configurable.

VERDICT: Thumbs up!

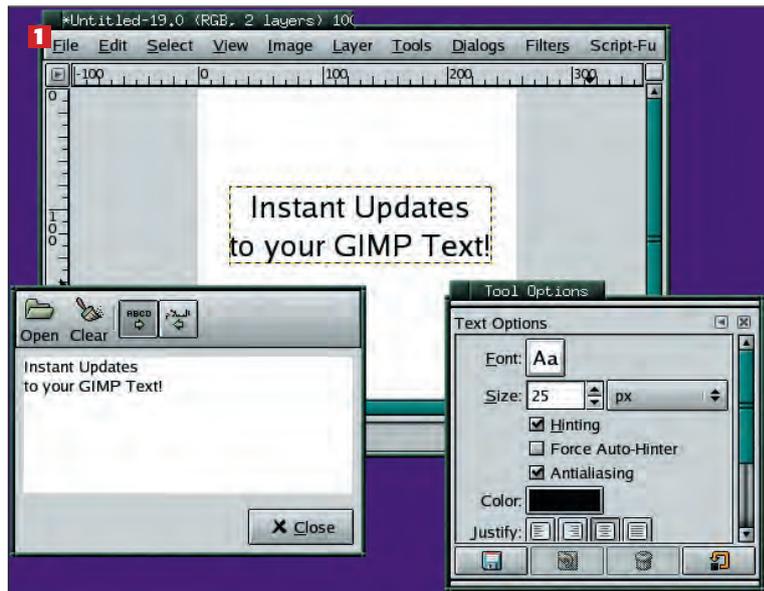
So did they get it right? Absolutely. The new layout is user friendly, easy to navigate and clutter-manageable. Menu options can still be set to user defined keystrokes, so there is no loss of functionality there. The highly polished icons in the Toolbox and dialogs make *The GIMP* look more like its proprietary cousins on Mac and Windows systems. Adding the 'Select by Color' option to the Toolbox is a big help too. In the 1.2 version, this very useful feature got buried in the Canvas Menu's Select sub-menu – and didn't have a Keyboard short cut or Keyboard accelerator. All these new features lead us to say the *GIMP* developers wiped the ease of use and clutter issues right off the canvas. ➤

TUTORIAL GIMP

WRITING IS IMAGERY

Another big failing of the 1.2 version of *The GIMP* was text handling. The built-in Text Tool was primitive by today's standards, and the problem was made only more complex with the addition of a second tool, better attuned to multiline text, but which was

not set as the default text tool. Font previews were unavailable and editing text was only possible if the secondary tool were used to create the text – the default tool didn't allow for text editing once the text was rendered.



Text Tool

1 *The GIMP 2.0* has but one Text Tool, but its tool options provide as much functionality as its predecessors and both the options and text editing window are easier to use. Text editing is done in a small preview window and changes are reflected immediately in the Canvas window. Editing is done by selecting the Text layer – which is now more easily identifiable by a Text icon in the Layers dialog. Multiline text is possible, including proper handling of newlines.

One of the other difficulties with using text in *The GIMP 1.2* was a lack of font previews. Choosing a font was simple enough – in fact, it might have been easier in 1.2 if you knew about font families or styles. But there has never been a facility to preview fonts as they would be rendered in a Canvas window. In *The GIMP 2.0* this issue has been addressed in the Text Tool Options dialog. The Font Selection button shows a sample of the currently selected font. Opening the Font Selection tool provides previews of all fonts in either list or grid views. Other options include setting the text justification and hinting, line spacing and indent.

The GIMP Text tool is greatly simplified, yet more flexible than its 1.2 counterpart.

VERDICT: Thumbs up, mostly

So did they get it right? For the most part, yes. From the ease-of-use perspective, the Text Tool is quite an improvement. Editing text is far easier and intuitive, and the results are immediately apparent. The unification of the best of both of the old text tools gives the new tool better flexibility and ease of configuration. Text rendering was fair in *The GIMP 1.2* and has improved with integration of the FreeType libraries in *The GIMP 2.0*. Overall, handling of text has greatly improved.

But all is not perfect here. Font previews are valuable, but resource-consuming if you have hundreds (or thousands) of fonts. The previews can definitely slow your system. And there is

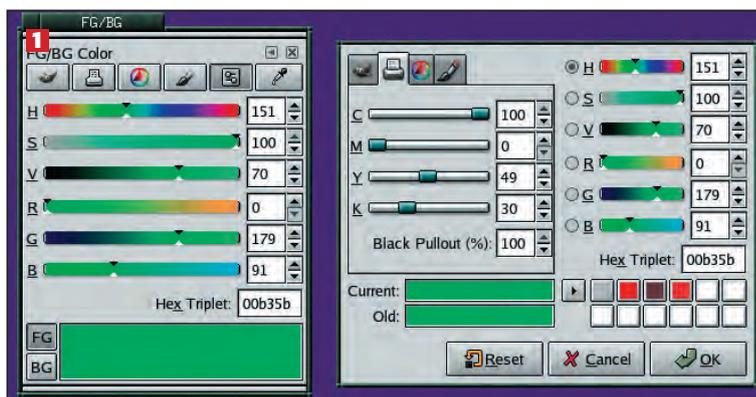
no configuration option to turn off the previews, so every time you have to select a font you can be slowed considerably.

Other issues are less obnoxious. The currently selected font name is not listed anywhere except in the font selection dialog. And while the direction of the font can be selected – an interesting option indeed – this option is in the text-editing dialog instead of the Tool Options. Considering that text justification is in the tool options, you'd think that other text characteristics would be there as well. Kerning is also not supported in *The GIMP 2.0*, even though the underlying FreeType libraries do support it.

Most of this is nitpicking, though. All things considered, we have to give a conditional thumbs up to the text handling improvements.

COLOUR MANAGEMENT AND OS SUPPORT

1 *The GIMP* is the tool of choice for Web developers working on Unix systems, but to go beyond this, the application needed to support multiple operating systems (besides Unix) and provide better colour management. The latter issue is a big one if *The GIMP* is ever to move seriously into the prepress arena.



***The GIMP* supports multiple platforms natively, but is that enough to keep users of those OSes happy?**

COLOUR MANAGEMENT AND OS SUPPORT CONTINUED

Colour Tools

1 In 2.0, the *GIMP* developers tackled multiple platform support the old-fashioned way – they let someone else do the work. *GTK2*, the software that gives the GNOME desktop its look, provides applications the foundation for running on multiple platforms with limited porting work. That geek-speak means that *GTK2* lets *The GIMP* run on Windows and Mac as a native application. While the *GIMP* developers still had plenty of work to do to clean the application for consistency across the three major platforms, most of the heavy work is handled by *GTK2*. The result is an expanded audience for *The GIMP 2.0*.

Colour management improvements came in bits and pieces with *The GIMP 2.0*. CMYK support – one of the most requested features for 2.0 – can be found in the Color selection dialog in the new Printer colour tab. This will be helpful when you know ahead of time the colour you will be using for a print job.

Other colour management improvements include a new colour mixer that allows selection of a colour from any part of the screen, not just from Canvas Windows. The Levels dialog now provides eyedropper selection of black, gray and white points. And the histogram is available as a dock, making it visible at the same time that the Curves dialog is being used.

VERDICT: Thumbs at half-mast

So did they get it right? Not really. OS support is a winner. More users these days are coming from the ranks of the unwashed Windows world than that of Unix. Mac use is growing fast, with OS X's support for X applications running side-by-side with native Mac applications. Having *The GIMP* on all three platforms makes the Unix bias less than it once was and, at a minimum, puts the application in front of the eyes of those who might not otherwise have seen it.

Colour Management, on the other hand, falls short in 2.0. The promise – implied or literal – was for 16-bit colour support in 2.0, and sadly we're seemingly no closer to that goal now than in December 2002. While there are many technical reasons for this, the problem is one of perception – the development team let the idea spread that the support was coming, specifically with 2.0. If wasn't going to happen, they should have nailed the lid down on that idea far sooner. Without 16 bit support, there is no chance of real colour profile support or colour gamut handling. 16-bit support is paramount for real-world prepress and video work. This issue aside, the improvements related to colour management are fairly minimal from the user perspective. The user interface improvements are not enough to win over the newly introduced Windows and Mac users, though: there has to be more meat on the bone.

The trouble here amounts to a PR failure more than a technical one, but no matter the reason, we have to give a thumbs down on colour management and platform improvements for 2.0. Plans call for 16-bit support (via another project called *GEGL*) in the next year; but that may be too late for the new crowds, and we've heard promises before... **LXF**

MORE GIMP-ROVEMENTS

More power, more options...

There is nowhere near enough space here to cover all the changes in *GIMP 2.0*. Some of the other notable improvements include:

- Improved vector (re: path) support.
- A visual undo history – you can see the changes previous edits created.
- Templates, such as PAL/NTSC, A4, Letter, and CD Cover, for new canvas windows.
- A visible grid that complements the use of guides.
- Built-in Session management.
- New layer modes: hard and soft light, grain merge and grain extract.
- Native Python support.

One change you won't see – or probably even notice – is an improved code base that will allow the developers to add features or extend existing features more easily than the code from 1.2. In fact, part of the reason for the delay in the 2.0 release was to clean the code up and make it easier to add features like 16-bit support. Time will tell if this hidden improvement will have been worth the wait. You'd have to be a real stick in the mud not to give the latest *GIMP* a hearty thumbs up. Developers may lack an understanding of PR, but that can't detract from obviously talented work. *The GIMP 2.0* offers most users better means to be productive and creative. And in the end, isn't that what we really want from an Open Source project?



The GIMP community is an integral part of the app's success, and there are some great online resources:
The GIMP: www.gimp.org, **GIMP Registry:** <http://registry.gimp.org/index.jsp>, **Graphics Muse Tools CD:** www.graphics-muse.com/gfxmuse/gfxmuse.html, **Tutorials and other useful links:** www.graphics-muse.com/cgi/gmcat.pl?id=11

MAKING MOVIES

Fading with Kino

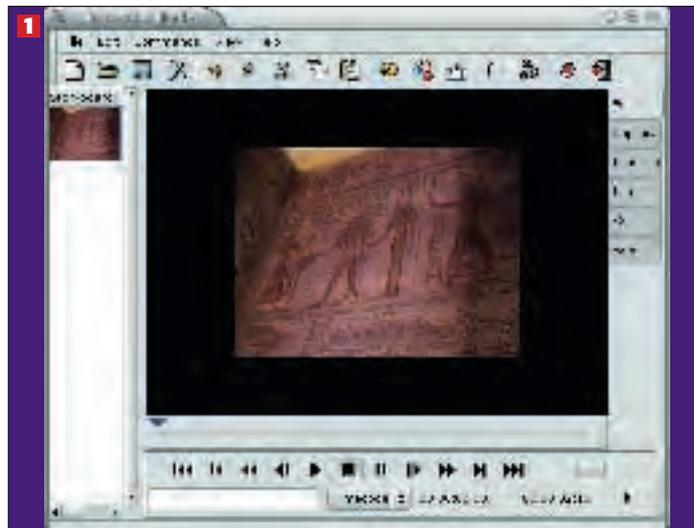


Biagio Lucini shows us how to add effects when we join our clips.

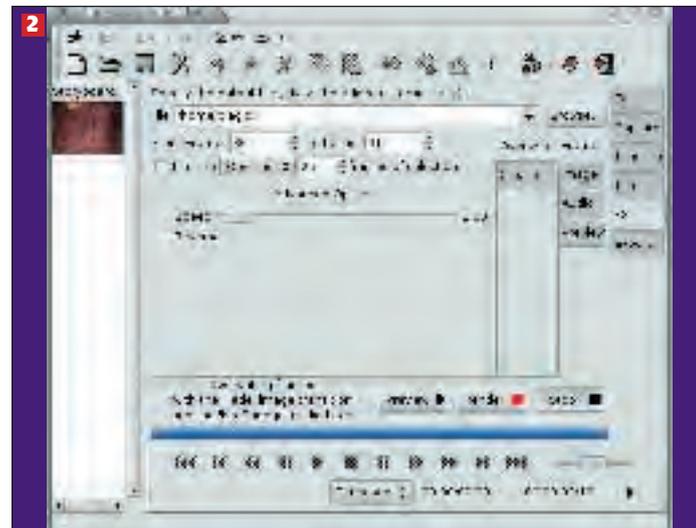
As promised in last month's cover feature, *LXF* would teach you some more about *Kino* and video editing in general in an occasional series: and here's the first one.

One of the most frequent operations in home video editing is gluing together two scenes with a transition effect in the

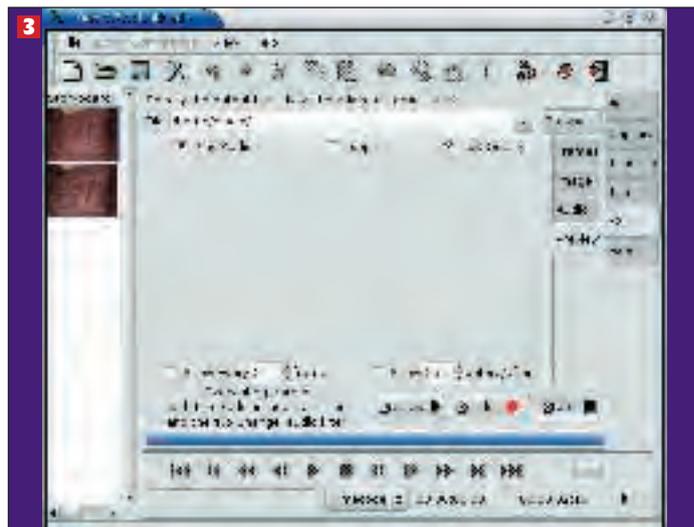
middle. Here, we will start with two clips we want to connect with a fade effect and highlight the necessary steps to obtain the desired result. *Kino* is one of the most flexible video editors for Linux, which has been described quite in detail in *LXF52*, to which you should refer for more general information about video editing in general.



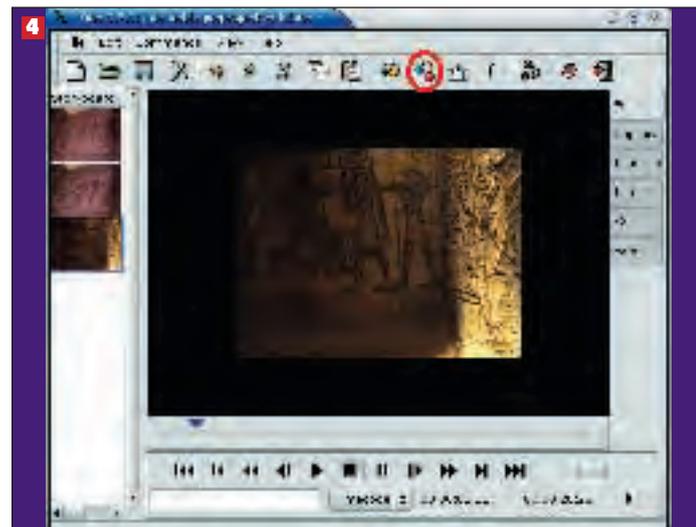
1 We start by importing the first clip: by clicking on File>Open, a browsing window pops up, in which the name of the file can be inserted. Now the clip should appear in the Storyboard pane on the left of your screen, as it does here in our example.



2 Click on the effect tab 'FX'. Next, select the tabs 'Frames' and 'Overwrite', and enter the numbers corresponding to the frames of our movie we want to overwrite with the effect in the first two boxes, remembering that 25 frames corresponds to around one second of 'movie-time'. For instance, in our example here, the clip has 61 frames. For a fade effect in the final part of the movie that lasts around one second, we enter the numbers 36 in the box 'From Frame' and 61 in the box 'To Frame'.



3 In FX>Image>Transition, select the effect 'Fade', then select 'Forward' from the drop-down menu in the lower left corner. The 'Preview' button gives an idea of how the clip will look with the selected settings – you can experiment a bit here by changing effect and duration. When ready, click the 'Render' button. Note how the scene has been split into two parts: look at how it is shown on the Storyboard – the first one is the unmodified part of the original clip, while the second, lower part has the effect in it.



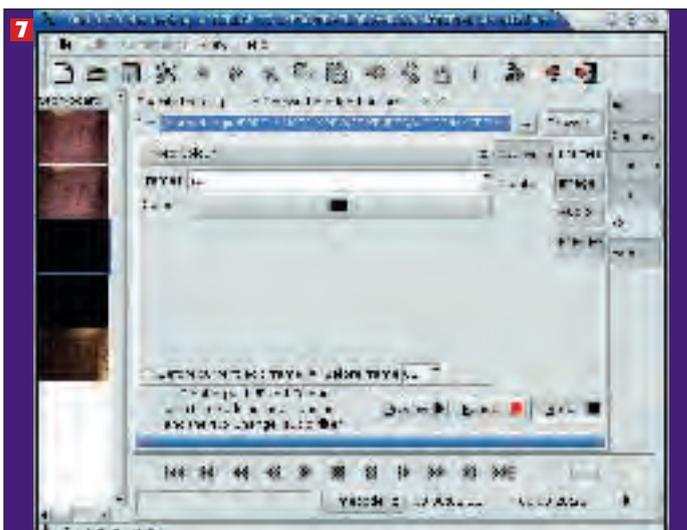
4 Now, go back into the Edit tab and select the end of the movie by using the movie navigation functions at the bottom of the view window. From the toolbar, click on 'Append File After Frame' (circled in red in the example above) and insert the name of the second clip. Now the second movie appears as the third scene from top to bottom in the Storyboard.



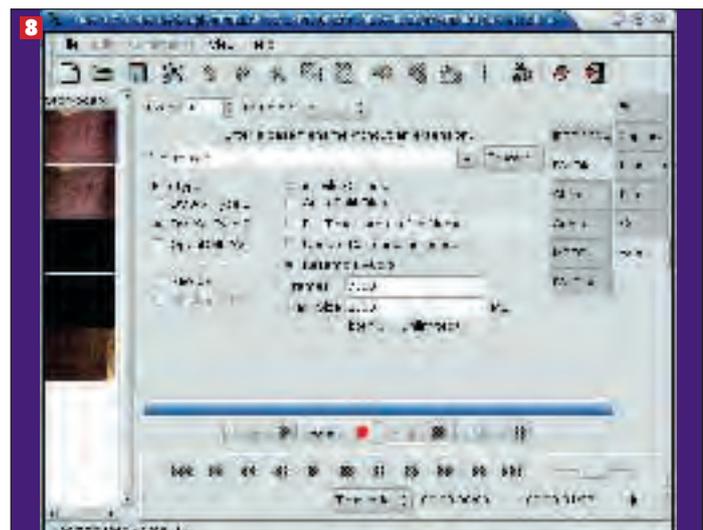
5 Select the file we have just inserted by clicking on the corresponding image in the Storyboard. We again go to **FX>Frames>Overwrite**. Now the default numbers in the boxes 'From Frame' and 'To Frame' will correspond to the first and last frame of the second clip. Leave the first value unchanged, then in the second field you need to insert a number corresponding to 50 frames forward. The effect in this part of the movie will last around 2 seconds.



6 Select **FX>Image>Transition**. Leaving the effect unchanged (Fade), we then select 'Reverse' from the first drop-down menu at the bottom and 'Frame Following' from the second. Then, check the box at the right of 'Frame Following'. You can click now on 'Render'. After this operation, you will notice a black image appearing in the Storyboard: this is the first scene of the freshly rendered effect.



7 We can now increase the duration of the effect for a smoother transition between the two clips. After having selected the black frame, go to **FX>Frames>Create**, and in the 'Frame' box, insert the number 75. By pressing 'Render', you will create 75 black frames (around 3 seconds of movie-time) before the start of the second clip. At last, it's time to actually render the movie.



8 Go to the tab **Export>DVfile**, and in the first line choose to export every single frame of the whole movie ('Every 1 frame of: All'). Then insert the name of the file that will contain the rendered movie, select DV2 as the export format and press 'Export', checking that the 'Autosplit' box is not selected. The operation is now complete and the final result can be admired with *MPlayer* or other similar video application. [LXF](#)

NETWORK ADMINISTRATION

Practical PHP Programming

The simple network management tools bundled with PHP mean you can centralise your administration – Paul Hudson shows you how...

Thanks to the efforts of our networking expert Dr Chris Brown in his dedicated SNMP tutorials last issue and this issue, I'm pleased to find most of the explanatory work about this protocol has been done for me already. That is, I'm pleased as much for your sake as I am for mine – although the Simple Network Management Protocol sounds like a walk in the park, the word "Simple" sits uneasily with its true nature. SNMP is not all that simple. In fact, if you have been reading Dr Brown's tutorials, you'll have been whisked on a whirlwind tour through its various uses, and should by now agree that we're all lucky there's no Hard Network Management Protocol because SNMP is hard enough!

If you haven't been reading Dr Brown's tutorials, I'd recommend you stop now and read them first. Rather than rehash everything he has said so far, I'll be assuming you're fairly fluent with everything he has covered. Expect to see regular references to OID (object identifier), MIB (Management Information Base), and other pieces of SNMP jargon. If you really have no opportunity to read the back issues, see the *Two-minute Summary* box on the opposite page.

CROSS-PLATFORM WARNING!

Windows SNMP concerns

Although PHP is largely cross-platform in that 95 per cent of scripts work as well on Linux as they do on Windows, utilising the SNMP functions will place your scripts into the elusive 5 per cent bracket. PHP's SNMP extension does

not work on Windows 95/98/Me because it relies on the NET-SNMP package, which is only available for Unix, and the Windows SNMP library is only available for Microsoft Windows NT/2000/XP/2003.

PHP and SNMP

Although SNMP supports data being sent when requested (polling) and data being sent when the device has encountered a set condition (notifications), PHP does not support both of these. Despite the CLI SAPI really getting into its stride back in the 4.3 release, this paradigm shift has yet to make it to the SNMP extension: the PHP SNMP extension is designed to contact SNMP devices and retrieve information, as opposed to the SNMP devices contacting PHP. Naturally this limits what you can do with PHP and SNMP, but only a little – we'll be looking at how you can make your PHP scripts connect to SNMP devices and retrieve status information. PHP has nominal support for SNMPv3, but don't expect too much – the extension is somewhat adrift in the PHP source tree, and doesn't get updated much beyond the occasional bug fix.

There are seven SNMP functions available, of which five communicate directly with your SNMP device. These seven are:

- **snmp_get_quick_print()** Gets the internal **quick_print** value. Takes no parameters, and returns true or false.
- **snmp_set_quick_print()** Sets the internal **quick_print** value. Takes a Boolean as its only parameter, and returns nothing.
- **snmpget()** Retrieve an SNMP object. Takes a minimum of three values: the hostname, community, and object ID. Returns the data you requested as a string.
- **snmprealwalk()** Retrieve all objects as an array with the OIDs as keys in the array and their values as the array values. This takes the same parameters as **snmpget()**.
- **snmpset()** If you have write access, this changes the value of an OID. This takes the same parameters as **snmpget()**, with

the addition of a variable type to set and and the variable itself.

- **snmpwalk()** This is essentially the same as **snmprealwalk()**, except that incrementing integers are used as the array keys rather than object IDs – this makes it less useful.
- **snmpwalkoid()** This is identical to **snmprealwalk()**.

The first two are the internal functions – by default, NET-SNMP returns its values with the data type prepended, eg

```
STRING: Foo bar baz
```

As this isn't helpful if you want to print it out without caring what the data type is, you can enable **quick_print** mode by using **snmp_set_quick_print(true)**. The default is to have **quick_print** disabled.

The two most important SNMP functions are **snmprealwalk()** and **snmpget()**: the former gets all values within a given object ID, and the latter just gets the value of the requested object. We'll use these two exclusively, with the onus on the former, for this tutorial.

First steps

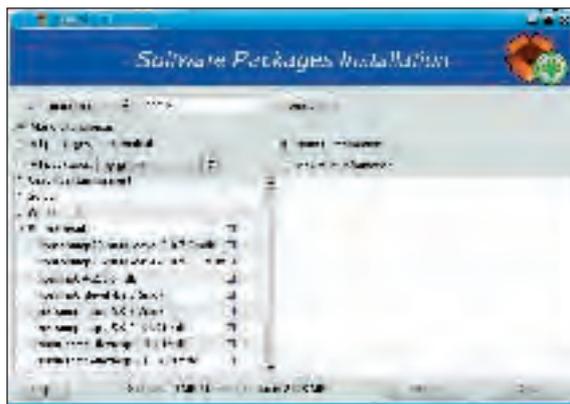
The easiest way to get started is to query your own machine for data, so we're going to get started using **snmprealwalk()**. Make sure you have your machine configured for SNMP with open access to the public community otherwise you'll need to use another host such as a router. Save this next example as `snmp1.php`:

```
<?php
$host = "localhost";
$community = "public";

snmp_set_quick_print(true);
$output = snmprealwalk($host, $community, "SNMPv2-
MIB::system");
var_dump($output);
?>
```

If you have SNMP set up correctly, that should output a lot of data: here's a cut-down copy of my output to give you an idea:

```
["SNMPv2-MIB::sysDescr.0"]=>
string(72) "Linux lxf-hudzilla 2.4.21-0.13mdksmp #1 SMP Fri
Mar 14 13:41:18 EST 2003 i686"
["SNMPv2-MIB::sysObjectID.0"]=>
string(33) "NET-SNMP-MIB::netSnmpAgentOIDs.10"
["SNMPv2-MIB::sysUpTime.0"]=>
string(12) "0:1:01:37.92"
["SNMPv2-MIB::sysContact.0"]=>
string(59) "Root <root@localhost>"
["SNMPv2-MIB::sysName.0"]=>
string(7) "lxf-hud"
```



Installing SNMP is easy if you have a friendly distro – just make sure you install the *-devel* libraries and the *utils* too!

SNMP – A QUICK RECAP

The two-minute summary

SNMP is an easy way to remotely request and receive information about computing devices on your network. On routers, this is useful for providing information about what levels of data have gone through on the network; on standard computers, you can query system name, contact information, uptime and more. As SNMP is so widely used – from switches to supercomputers – you probably already have some devices in your house that can be accessed through the protocol. If you don't own a router, switch, or particularly advanced fridge, never fear – your computer will do just fine. Note that SNMP is not necessarily limited to TCP/IP networks.

As Internet protocols go, SNMP has been around a good while. The earliest official document describing SNMP is RFC1065, which

dates from 1988. SNMP has remained under active development and there have been two new major releases, version 2 (1993) and version 3 which emerged around the turn of the century. RFC (Request for Comments) documents are the 'official' way to propose new Internet protocols. Their issue is regulated by the IETF (Internet Engineering Task Force) and over time some of them become adopted as formal Internet standards. There are many RFCs describing various aspects of SNMP; not all RFCs aspire to become full standards. Some are of a more informational nature. One of Dr Brown's favourites is RFC 1118, entitled *Hitch-Hiker's Guide to the Internet*. RFCs are available from many Internet archives, but a good place to start is www.ietf.org/rfc.html

```
["SNMPv2-MIB::sysORID.1"]=>
string(13) "IF-MIB::ifMIB"
["SNMPv2-MIB::sysORID.2"]=>
string(19) "SNMPv2-MIB::snmpMIB"
["SNMPv2-MIB::sysORID.3"]=>
string(15) "TCP-MIB::tcpMIB"
["SNMPv2-MIB::sysORID.4"]=>
string(10) "IP-MIB::ip"
["SNMPv2-MIB::sysORID.5"]=>
string(15) "UDP-MIB::udpMIB"
["SNMPv2-MIB::sysORID.6"]=>
string(39) "SNMP-VIEW-BASED-ACM-MIB::vacmBasicGroup"
["SNMPv2-MIB::sysORID.7"]=>
string(46) "SNMP-FRAMEWORK-
MIB::snmpFrameworkMIBCompliance"
["SNMPv2-MIB::sysORID.8"]=>
string(31) "SNMP-MPD-MIB::snmpMPDCompliance"
["SNMPv2-MIB::sysORID.9"]=>
string(40) "SNMP-USER-BASED-SM-
MIB::usmMIBCompliance"
```

As you can see, each array key there is a qualified SNMP object ID, which means you can query each of them individually by using those names if you want to – you will probably get quite a few more on your system, because the script outputs all objects under the system root. Note that the object IDs, eg **SNMPv2-MIB::sysName.0**, need to be used exactly as shown – they are case-sensitive, punctuation-sensitive, and just all-round sensitive so do take the output of **snmprealwalk()** seriously.

Note that each of the variables has no type information. In Dr Brown's SNMP tutorial in *LXF52*, he discussed the various data types used in SNMP – counters, gauges, enumerations, timeticks, strings, etc – and these are usually prepended onto the values as they are sent. This is the effect of **quick_print**: if it were set to false, either by default or through an explicit call to **snmp_set_quick_print(false)**, each of the values would have been output like this:

```
["SNMPv2-MIB::sysDescr.0"]=>
string(80) "STRING: Linux lxf-hudzilla 2.4.21-0.13mdksmp #1
SMP Fri Mar 14 13:41:18 EST 2003 i686"
```

It's quite helpful to have the data type information in there so you can distinguish between a count value and a gauge, both of which appear the same, but if you all you want to do is print out system information without parsing it all, leave **quick_print** on. >>

TUTORIAL PHP

◀◀ With a little tweaking to the original script, you can make the output a little easier to understand (and easier on the eye). Save this next code block as `snmp2.php`:

```
<?php
$host = "localhost";
$community = "public";

snmp_set_quick_print(true);
$output = snmprealwalk($host, $community, "SNMPv2-
MIB::system");
foreach($output as $snmpoid => $snmpvalue) {
    echo "$snmpoid: $snmpvalue\n";
}
?>
```

This time you should get the output formatted neatly, like this:

```
SNMPv2-MIB::sysContact.0: Root <root@localhost>
SNMPv2-MIB::sysName.0: hud-lxf
```

Unless you plan on creating a 'system status' page outputting all available data about a device, using `snmprealwalk()` is not helpful for much more than just seeing what object IDs are available. Even then, if you have installed the NET-SNMP utils package (see the box, *Installing SNMP support* for instructions), you can use the `snmpwalk` tool to accomplish the same task, like this:

```
snmpwalk -v 1 -c public localhost SNMPv2-MIB::system
```

If all you want is to query just one value, we need to get onto the `snmpget()` function. Note that if you want to see what other management information is available outside the 'system' root, just specify "" as the third parameter – this will retrieve *all* information available on the device, and may well take several minutes to run. If you are going to do this, I suggest you pipe the output to a file so that it doesn't overflow your terminal buffer. There are several other tools bundled with the NET-SNMP package as well as `snmpwalk` – try typing `snmp` then hitting your Tab key to see the available commands. You'll probably want to try out both `snmpget` and `snmpset`.

Getting single values

If you're getting three or more values at once, it's more economical to use `snmprealwalk()` to retrieve all the values within a group then extract what you're interested in, but if you just want one or two values it's easiest to grab them with `snmpget()`.

This takes the same basic three parameters as `snmprealwalk()`: the host to connect to, the community, and object ID to read, and returns the value of the object. Also as with `snmprealwalk()`, you should enable `quick_print` by calling `snmp_set_quick_print(true)` before using it so that you don't get the data types printed out with values. The key difference

between `snmprealwalk()` and `snmpget()` is that the third parameter in the first instance returns is used as the root, which means that all objects beneath it are returned, whereas in the second instance it is the absolute value to fetch, so it is returned immediately. What this means is that the keys set in the array return value from `snmprealwalk()` can be used as the values to fetch in `snmpget()`.

Here's a chunk of code – save it as `snmp3.php`:

```
<?php
$host = "localhost";
$community = "public";
snmp_set_quick_print(true);
$val = snmpget($host, $community, "SNMPv2-
MIB::system.sysUpTime.0");
print "SNMPv2-MIB::system.sysUpTime.0 is: $val\n";
?>
```

Looks largely the same as the last script, right? Well, yes – once you have a grasp of the basics, you can chop and change your actual SNMP calls fairly easily. Note that `sysUpTime.0` returns how long the SNMP service has been up – if you just installed your SNMP service for this feature, this will not reflect your actual system uptime.

Setting a value

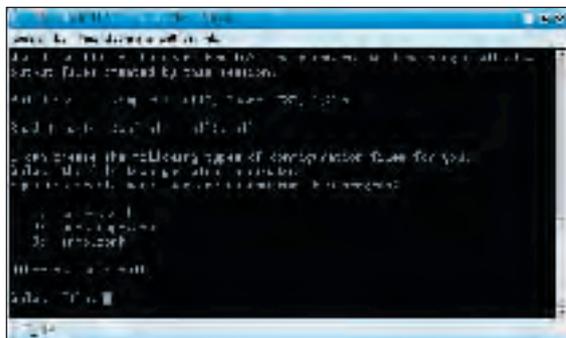
The complement function to `snmpget()` is of course `snmpset()`, and is only a little more complicated. As well as providing the host, community, and object to work with, you also need to provide a value type to set and the value itself. The type to set (parameter four) must be a single character from one of the following:

I	Integer
t	Timeticks
a	IP address
o	Object ID
s	String
x	Hexadecimal string
d	Decimal string

Finally, the value you're setting comes in parameter five. Note: Before you change any values on your server, you need to have a community in place that has read/write access to the data. Once that is organised, you can start writing editing scripts like this:

```
<?php
$host = "localhost";
$community = "private";
snmp_set_quick_print(true);
$val = snmpget($host, $community, "SNMPv2-
MIB::system.sysName.0");
print "SNMPv2-MIB::system.sysName is: $val\n";
$success = snmpset($host, $community, "SNMPv2-
MIB::system.sysName.0");
```

You can configure your local SNMP server through the `snmpconf` program, which guides you question-by-question through basic configuration options.



PROBLEMS WITH ICC?

Seg faults and such...

Although it's hard to draw any conclusions, I had serious problems with the PHP SNMP extension when compiled with v8 of Intel's *C++ Compiler* – the CLI SAPI would randomly segfault, nearly always when `snmpget()` was called. This problem disappeared entirely when I switched to *GCC*, which makes me think there's some sort of miscompile. If you don't use *GCC* as your compiler, be wary – if your scripts terminate abruptly with "Segmentation fault", you may have a bad binary.

```

if ($success) {
    $val = snmpget($host, $community, "SNMPv2-
MIB::system.sysName.0");
    print "SNMPv2-MIB::system.sysName is now: $val\n";
} else {
    print "Set of SNMPv2-MIB::system.sysName.0 failed -
check your SNMP access control settings!\n";
}
?>

```

Setting values in SNMP is often equivalent to changing status – setting a 'FooEnabled' value to 0 usually acts to disable Foo, so be careful what you change.

Conclusion

Although SNMP isn't a simple protocol by any means, PHP has limited or non-existent support for much of the advanced functionality – this is probably irritating for more advanced users, but certainly makes learning easier! As there are only four functions you'll need to use regularly, of which three are actually used in any complex ways, it should take you just a few minutes to be proficient in its use, which means you can spend more time using SNMP and less time learning about it.

If you want to dig more into the depths of SNMP and learn how to do more complex operations with it, read Dr Brown's SNMP tutorials in *LXF52*, *LXF53* and *LXF54*. Sadly, there are no PHP books available at this time that give sufficient coverage to SNMP for me to want to recommend them, and even the manual is a little bit sparse. Technically, this places you at the frontier of

INSTALLING SNMP SUPPORT

Find the relevant libraries

To enable support for the SNMP extension in your PHP build, you first need to download and install the NET-SNMP libraries from www.net-snmp.org/download. Many distros, including Mandrake, come with NET-SNMP packages ready for use. Once you have it installed, go into your PHP source directory and run configure again, this time adding `--with-snmp` and `--enable-ucd-snmp-hack` to your previous configuration.

If you compiled PHP previously, it will only rebuild the SNMP extension, so this should only take a moment. If you get errors about PHP being unable to find `snmp.h`, you haven't fully installed the NET-SNMP library – this is likely to

be in a different package name in your distro. Depending on how smart your distro is, you may need to install the `-devel` libraries for `libelf` and `libnet-snmp` also. Errors such as "Cannot find module (IP-MIB)" means that your distro has opted not to automatically install the `net-snmp-mibs` package – you'll need that too.

Finally, to make life easy, make sure your distro installs the `net-snmp-utils` package – this includes tools such as `snmpwalk`, which make it a great deal easier to debug your SNMP server.

If you're short on time, just search for 'snmp' in the available packages and install all the latest one versions available – it's not ideal, but it should work first time!

programming, which is quite fun once you get used to it – get out there and try some new things out to see how SNMP can give your site a boost. Yes, half the things you try will be too crazy to work, but that's all part of the fun!

One last note: please be very careful with allowing write access to your SNMP device. By default, the public community usually has read access or sometimes even no access at all, so it's best to create a new community specifically for writing. Even then, you should be careful about what data you allow to be changed – it's very easy to do some quite serious damage to an overly trusting SNMP server, and it's best that you play the "better safe than sorry" card. [LXF](#)

NEXT MONTH

Next issue we'll either be looking at the **POSIX functions** or the **reflection API**, depending on how the wind blows – if there's space, you might even get both!

unique customers - unique hosting

A FREE MONTH
of web hosting.
When ordering online
use "FUTLIN053".

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/ASP/XML

from **£7.95 pcm**

Gold

- 200Mb Disk Space
- 4Gb Bandwidth
- 10 POP3 e-mails
- PHP/ASP/XML
- 1 MySQL / Access DB

£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 0808 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

DATABASE ADMINISTRATION

Server School: SQL database servers

Flat files not fast enough? Roger Burton West examines SQL servers.

When storing data in flat files, be they plain text or XML, has stopped being fast enough for the job, the obvious next step is an SQL database.

Although database servers have traditionally been heavyweight processes, a modern system's footprint is very much proportional to the demands placed on it, making an SQL server a reasonable choice even for relatively small projects. This month's tutorial deals with what the server administrator needs to know, with diversions into database choice, ways of assuring data integrity, and some of the features that separate the major implementations.

SQL is a standardised language for issuing database queries and commands. Unfortunately, while the basics are indeed much the same from one database to another, the extensions are not, and most of the database servers have independent reinventions of common functions. There are many books and websites that introduce the basics of SQL; it is worth getting hold of one that is relatively neutral in dialect and reading it in combination with the database manual rather than using something vendor-specific. The extra time investment is minimal, and when (as will inevitably happen) you need to port code from one database to another a general familiarity is helpful.

Getting started

The main players in the Free database field today are *MySQL* and *PostgreSQL*. Both of them are long-term projects, and both have their zealots and detractors. As a very general guide, *PostgreSQL*

started off as a research project and has always been a very full-featured database, but is sometimes compromised on speed; *MySQL* began as a system to handle databases with high speed, and many important database features (in particular, transactions) have only been added in a piecemeal way in recent versions.

Generally speaking, *MySQL* and *PostgreSQL* are roughly equal in speed, though in specific applications one or the other will typically have a slight advantage. If speed is absolutely crucial, it is worth trying both on the sort of data you will be using.

Installation of a database server is slightly more complex than a standard package installation, because you will normally need to set up administrative passwords for the system. As a rule of thumb, these should only ever be used for high-level server administration, not for work on individual databases; both *MySQL* and *PostgreSQL* allow the creation of additional users with restricted privileges, much as a well-run Unix system doesn't execute every process as root.

Each server comes with its own command-line tool for entering SQL statements and commands; there are also add-on programs available to assist in server administration. Almost all programming languages have libraries available for database access.

Creating a database in *MySQL* is simple: just **CREATE DATABASE accounts** and **GRANT ALL** to the owner of the database. This will also work under *PostgreSQL*, but because of the way the ownership system works it can cause some problems. Ideally, the user who is going to be administering the database should be its creator (and thus its owner), but this user

SQLITE

The minimalist approach

SQLite (www.sqlite.org/) is a database with a difference: there is no separate server process, and everything is done by file manipulation (usually on the local machine). It is extremely fast (roughly twice as fast as *MySQL*), supports full transactions (though very limited constraints and triggers, and no separate stored procedures), and has bindings/drivers available for most languages (including Perl, C, Ruby, Python, PHP and Java).

Where it falls down is on concurrent updates. The engine operates with a global file lock, so any update will block other updates until it has been completed. *SQLite* is not well-suited to a system that has multiple updates going on simultaneously. For a straight content server in small and medium-sized applications, though, there is little faster.

TIP: The slowest part of *SQLite* is writing to the filesystem. By making your transactions relatively large (for example, when adding to a database, only commit every thousand rows or so), you will find that it runs much faster.

may well not have database creation permission; and there is at present no way to change the ownership of a database. The best way round this with current versions of *PostgreSQL* is a temporary grant of the **CREATEDB** privilege by the **postmaster**:

```
(postmaster) ALTER USER bob WITH CREATEDB;
```

```
(bob) CREATE DATABASE accounts;
```

```
(postmaster) ALTER USER BOB WITH NOCREATEDB;
```

Security

The short version of database security is: no database server should ever be exposed to a public network. Although these days they are mostly quite robust, both *MySQL* and *PostgreSQL* have had significant remote-access vulnerabilities (compromising either the database or the entire server on which it is hosted), and there is no guarantee that this will not be repeated; moreover, it is rarely necessary to grant the world access to a database server, in contrast to a web or mail server. So the first stage in securing a database server is to firewall it on the relevant ports (by default, 3306 for *MySQL*, 5432 for *PostgreSQL*) from all hosts except the ones that actually need to access it. This is not so much an alternative to the access restrictions described here as a supplement to them: a strategy of defence in depth is as sound here as it is elsewhere.

In most cases, the typical database server will be hosted either on the same machine as the front-end programs that are needed for users to communicate with it, or nearby on the network. In either case, you will have a list of IP addresses that require access; the second layer of defence is configuration of the TCP/IP listener service.

If a *MySQL* server is only going to be accessed from the local machine (via Unix sockets), use the **--skip-networking** command-line parameter to *mysqld* to disable the TCP/IP listener. Similarly, the **-i** parameter to *PostgreSQL* will enable TCP/IP (by default it is not available). For a system with multiple interfaces, the **-h** option to the *PostgreSQL* **postmaster** will restrict the server to listening only on the IP address specified.

Both *MySQL* and *PostgreSQL* support access restriction by IP address, which is the third stage of securing the server. Under *MySQL*, this is combined with user management (and indeed is the preferred way of creating users):

```
GRANT privileges ON database.* TO user@'hostspec'
```

where **privileges** is the set of allowed actions (eg **SELECT** or



pgmonitor lets you keep track of what the *PostgreSQL* server is doing.

ALL PRIVILEGES), **database** is the name of the database which you wish to affect (.* meaning that it will take effect on all tables within that database), **user** is the *MySQL* username to be changed, and **hostspec** defines the host or hosts which are allowed access. This can be a host name (checked against the reverse DNS name of an incoming connection), IP address, or IP with dotted-quad netmask. In each case, % is a wildcard for a section – so % on its own would allow access from anywhere, **%mycompany.com** would allow access to all **mycompany.com** machines, and **192.168.0.%** would be equivalent to **196.168.0.0/255.255.255.0**.

A username created in this way will have no password. Add the clause **IDENTIFIED BY 'password'** to the command to restrict access.

Under *PostgreSQL*, host access control is managed through the *pg_hba.conf* file. This is often one of the most troublesome areas of *PostgreSQL* for the new administrator, but can be dealt with quite easily step by step. First, it is necessary to create the relevant users:

```
CREATE USER bob PASSWORD 'password';
```

There are two basic forms for access lines in *pg_hba.conf*. To grant access to connections from the local machine:

```
local (database) (user) (authentication info)
```

And to grant remote access:

```
host (database) (user) (host-spec) (authentication info)
```

The initial **host** can be replaced by **hostssl** or **hostnossl** to allow only SSL-encrypted connections, or only unencrypted connections, respectively.

(database) specifies the individual database to which this user will have access. **all** allows access to all databases.

samegroup allows access to all databases for which the user is in the group named identically to that database (see **CREATE GROUP** and **ALTER GROUP**). Multiple database names are separated with commas.

(user) is the username. **all** allows access to all users. Older versions of *PostgreSQL* do not support user specification.

(host-spec) is a host definition. Only IP addresses may be used, not DNS names, and both a network and a netmask must be given. This can be in either of the commonly used forms, so **192.168.0.0/24** and **192.168.0.0 255.255.255.0** will both work (note that there is no slash in the dotted-quad netmask form). To specify a single IP address, just give it a **/32** netmask.

(authentication info) gives the basic authentication method with optional extra parameters dependent on that method. This can be one of: **trust**, which allows unconditional access – useful for development work, not for production servers; **md5** (or **crypt** or **password**), which allow login with a password (in decreasing



◀◀ order of encryption strength) and are the methods most commonly used; and **krb4**, **krb5**, **ident** and **pam**, which allow logins with *Kerberos*, *identd* lookup, and pluggable authentication modules (very rarely used in production environments, but invaluable on the rare occasions when they are wanted). Finally, **reject** can be used to deny authentication – which is very handy if you want to allow most – but not all – of the machines in a subnet to connect.

If you need to connect to a database server across a public network – for example to administer a live server in a colocation facility – be aware that while both *MySQL* and *PostgreSQL* will make some effort to obscure the contents of passwords (*MySQL* from v4.10, *PostgreSQL* with the *md5* or *crypt* authentication methods), the remainder of the connection is vulnerable to packet capture. Rather than relying on this built-in cryptography, encrypt the entire TCP/IP connection using *SSH* or *SSL*; this also allows connection from arbitrary IP addresses without opening the server ports to the world, since the tunnelled connection will appear to originate from a local host.

Permissions

Typically there will be as many as three or four users associated with any given database, each with a different set of access permissions:

- 1 The 'root' user – **postmaster**, under *PostgreSQL*. This is the server administrator. Normally operated by a human or with administrative scripts. Only the server administrator should have this level of access.
- 2 The database's owner. This user has full administrative rights over the specific database on the server, and can create new tables, amend data, and perform other required operations. Normally operated by a human or with administrative scripts.
- 3 The 'update' user. Application programs needing to apply changes will connect as this user; they cannot, even if compromised, modify the database's table structure.
- 4 The 'read-only' user. This is also used for application programs, typically in the content-delivery role. If the program is compromised, the database will remain intact.

Where the application is structured so as to permit more finely-grained access control, these may be subdivided further (for example, a read-only user with rights to update a single table of access counters), but the model given here is a good starting point. One persistent problem with application programs relying on an SQL database is that many of them have been found to have bugs which allowed database access (for example, search the web for "SQL Injection Attack"); by their nature, these often have to be available to the outside world, but the security risk can be minimised by granting only the privileges absolutely necessary to a task. Obviously you hope to have a secure web application as well, but... defend in depth.

PostgreSQL defaults to a closed-access model; it will be necessary for the owner of the database to make explicit grants of privileges to the users that need them, on a per-table basis. For example:

```
GRANT SELECT ON users TO reader;
GRANT SELECT,INSERT,UPDATE,DELETE ON users TO updater;
```

With *MySQL* the syntax is very similar, but it is necessary to specify both the database and the table:

```
GRANT SELECT,INSERT,UPDATE,DELETE ON application.users
TO updater;
```

Performance

MySQL offers two major table types: the traditional and fast MyISAM, and the newer InnoDB. MyISAM does not support transactions, triggers, or many other features considered part of a standard database, but uses little memory and runs quickly; InnoDB has all these features but is significantly slower and requires more memory. It is entirely possible to create a database with mixed table types, but a transaction on mixed tables cannot (for obvious reasons) be rolled back if part of it fails; this can cause referential integrity problems even in the InnoDB parts of the database.

If you are contemplating a structure with mixed important-but-slow and expendable-but-fast, a safer approach is to move the MyISAM tables to a separate database (either on the *MySQL* server or a separate process entirely, such as *SQLite*); this enforces a true separation of the protected from the expendable data at the application level. It is better, of course, simply to use transaction-capable tables throughout.

Other table types in *MySQL* are BDB (transaction-capable, using a standardised *BerkeleyDB* back-end; quite slow but useful if you want to access the database with other tools); MERGE tables (a collection of MyISAM tables with identical column definitions, accessed as a group); and HEAP tables (stored only in memory, so very fast but ephemeral). *PostgreSQL* offers only transaction-safe tables in a single format.

Both *MySQL* and *PostgreSQL* will log every query made to the database engine. This is often useful during the debugging of a web application, but can be a problem when site traffic is high; an aggressive log rotation and compression policy will normally prevent this from causing difficulties. It is not recommended that logging be turned off unless daily log rotations are not sufficient, and even then only when all the application code to be used with the database is known to be very stable. Since both databases normally store data as well as log information on the /var partition (the exact location varies by Linux distribution), a database server for a large system will need slightly unusual partitioning to perform effectively; either the database root (as a separate partition) or /var itself can profit greatly from being on a RAID system or other fast-seeking hardware.

PostgreSQL is supplied with the **VACUUM** command, which performs garbage-collection and optimisation on a database. Many installations have this installed to run automatically overnight; it can take some time and system resources, but will improve performance significantly at other times. With the **ANALYZE** command (which can be run at the same time or separately), *PostgreSQL* will gather statistics to be used in its query optimiser.

Integrity

The true advantage of the SQL database is not merely in retrieval speed or convenience of access, but in the ability to maintain the integrity of the data and act automatically upon it. An SQL database field is capable of containing a great variety of data, of which only a small subset may be meaningful to one particular application; how can you make sure that bad data doesn't get into the database?

You could of course specify this in the documentation, and trust the application developer to get this right, check for errors every time the table is updated, and so on. And indeed you should; but since this does not give a 100 per cent guarantee, you can also cause the database itself to reject entries that would put it into an inconsistent state.

For a first example, consider a table of names:

```
CREATE TABLE user (
  id INTEGER,
  name VARCHAR(40)
);
```

Any record in this table must have a three-digit number ID, and a name that isn't blank. In *PostgreSQL* this is a simple constraint:

```
CREATE TABLE user (
  id INTEGER,
  name VARCHAR(40),
  CONSTRAINT con1 CHECK (id > 99 AND id < 1000 AND
  name <> '')
);
```

MySQL has a more restricted constraint checking set but can still verify uniqueness of index values.

A foreign key constraint provides slightly more complex example: a second table, **user_property**, will have a **user_id** field, which must refer to a user defined in the user table. The syntax here is the same in both *PostgreSQL* and *MySQL*:

```
CREATE TABLE user_property (
  user_id INTEGER,
  FOREIGN KEY (user_id)
  REFERENCES user(id)
);
```

though this is one of the many features of *MySQL* which is only usable with the relatively slow InnoDB table type.

Triggers generalise this idea further, allowing stored procedures – arbitrary programs in SQL or another language – to be run automatically on the server when data are added, changed or deleted. They tend to be very database-specific, as the syntax and even the language used by the procedures changes a great deal; but in a complex system it is normal for a significant part of the 'business logic' to be encoded here rather than in the CGI scripts or other systems placed closer to the outside world. This generally runs faster than the equivalent program in a general-purpose language; cuts down on network traffic and other sources of delay; makes it easier to code alternative front-end systems; and assists security, if an organisation denies users direct access to the database but requires that all changes go through stored procedures (for enhanced parameter validation and logging).

At present, *MySQL* supports stored procedures (from version 5.0, which at the time of writing is under continuing development but available in the *BitKeeper* version control system) but does not yet support triggers. It uses the SQL:2003 syntax (as seen in *DB2* from IBM). This is a very new feature, and is likely to change as development continues.

PostgreSQL has had stored procedures for rather longer; the simplest option is to write them in PL/pgSQL, which is a *PostgreSQL*-specific SQL variant, though Perl, Python and Tcl are also supported, and many people are still using the original C-language function extension system.

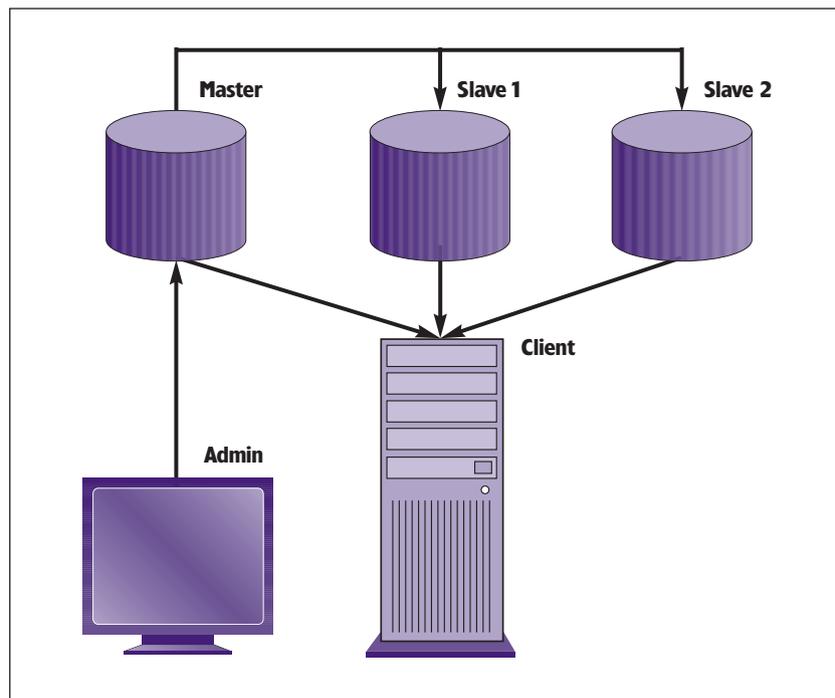
It is necessary to enable the use of a language within a specific database:

```
createlang plpgsql (database);
```

If **template1** is used as the database name, the language will be available in all databases subsequently created.

Backups and replication

The problem in backing up a database system is that the data may change while the backup is taking place. For this reason,



most databases offer a means of obtaining a 'snapshot', which may be used to recreate the database in the event of problems.

MySQL offers the **mysqldump** command, which will produce a consistent view of the data as a script file containing the SQL commands needed to recreate the database. It is normally run with the parameter **--opt**, which ensures that the dumps are as consistent as possible (though of course there may still be inconsistency between separate databases) and allows the fastest possible loading

PostgreSQL offers **pg_dump**, which does essentially the same job. It has two output options: a script file, as **mysqldump**, and an 'archive file' to be used by the **pg_restore** command. This latter allows selective restoration of the data without a need to edit the (possibly very large) dump file.

MySQL supports a one-way replication system: updates on a single master server can be automatically distributed to slave servers. This is not a truly redundant system or one that can help to share the load of updates, though it is certainly helpful where the main load on a system is database reads. The setup procedure essentially involves freezing the database on the master and copying it to the slave server, so this is best done before a system enters production.

Since *MySQL*'s replication depends on the updates being replayed as commands on the slave server, it is not quite as secure as a true replication mechanism. Several users seem to have experienced database corruption when network connections are problematical. This can of course always be fixed by manually resynchronising the databases.

PostgreSQL does not currently support replication; this feature is expected to be added in the near future.

Summary

While an SQL database needs a significant investment in both coder and administrator knowledge, the consequent performance gains that become available can be very significant. It is not a solution to every problem, but it makes a great many problems so much simpler to solve. **LXF**

MySQL replication only allows one-way data flows but is still useful for taking the load off the main server.

NEXT MONTH

What good are defences if you can't actually tell whether your server has been breached or not? Next issue we demystify the outs and ins of intrusion detection.

PERL TEMPLATE TOOLKIT

Configuration File Templates

PART 2: OPTIONS Dave Cross makes generating configuration files easy with the Template Toolkit.

Last month we were processing one template at a time using *tpage*. Often you will want to process a set of associated templates at the same time. The Template Toolkit comes with a utility program called *ttree* that allows you to do just that. It processes a set of files from an input directory and puts the processed versions into an output directory.

There are many other ways that *ttree* is more powerful than *tpage*. It supports a huge number of options that control exactly how your templates are processed. You can get a preview of what all of those options are by typing **ttree -h** or a more detailed description by typing **man ttree**. We'll cover some of them in more detail later.

Networking configuration files

In this article, we'll create some configuration files that define a simple network. Specifically, we'll generate an `/etc/hosts` file and some of the files that are required to configure *BIND*. We'll be looking at the simple network that which is described in the boxout *The Sunnysdale Network* on the opposite page.

Getting started with *ttree*

As I mentioned before, *ttree* takes templates from an input directory, processes them, and writes the processed versions to an output directory. Therefore when starting with a new *ttree* project, I like to create a project directory and subdirectories called **in** and **out**.

You can configure *ttree* in a number of ways. The easiest is probably to use a `.ttreerc` file. By default, *ttree* looks for `~/ttreerc`, but you can override this by either using the **-f** option or by setting the `$TTREERC` environment variable. As I often like to have slightly different *ttree* configurations for different projects, I set `$TTREERC` to be `./ttreerc` and put a new `.ttreerc` file in each project directory.

One nice little touch is that if you run *ttree* and you don't have a `.ttreerc` file (either in your home directory or in the location defined by `$TTREERC`) then *ttree* will offer to create a sample file for you. This file has most of the common *ttree* options in it together with copious comments that make it easy to edit. We won't use that this time, we'll edit our own `.ttreerc` from scratch. Here are the first three lines:

```
verbose
```

```
src = ./in
```

```
dest = ./out
```

The second and third lines are pretty self-explanatory. They define the source and destination template directories. The first line puts *ttree* into verbose mode where it tells you everything that it is doing.

With this in place we can test our first template. We'll create a template that expands our network files into a hosts file. Here's the template.

```
#
# /etc/hosts
#
[% USE networks = datafile('data/networks.txt') -%]
[% FOREACH network = networks -%]
# Network: [% network.netname %]
# IP: [% network.number %]
[% USE hosts = datafile("data/net_${network.netname}.txt") -%]
[% FOREACH host = hosts -%]
[% host.IP %] [% host.hostname %] [% host.alias %]
[% END -%]
# End of network [% network.netname %]

[% END -%]

# End of networks
```

If you put this in a file called 'hosts' in the `./in` directory, you can then run **ttree** like this

```
$ ttree
```

If all is well, you will see output that looks something like this:

```
ttree 2.75 (Template Toolkit version 2.13)
```

```
Source: ./in
Destination: ./out
Include Path: [ ]
Ignore: [ ]
Copy: [ ]
Accept: [ ]
Suffix: [ ]

+ hosts
```

ABOUT THE AUTHOR...

Dave Cross is a full-time Perl consultant and the author/co-author of a number of books on the subject, including the recently published *Perl Template Toolkit* (O'Reilly, ISBN: 0-596-00476-1)



You'll see that *ttree* has reported on the source and destination directories that it has processed. It also reports on a number of other options which we will explain soon. At the end of the output you'll see the name of our template file. The **+** sign next to it indicates that the template was processed.

ttree has laziness built-in. It will only process the templates that are necessary. It works out which templates to process by comparing the contents of the source and destination directories. A template is only processed if it doesn't exist in the destination directory or if the source version is more recent than the destination version. In fact it works very much like *make*. If you try to run **ttree** again immediately, then the output will be identical to the first run, except you will see that the last line will be replaced with:

```
- hosts (not modified)
```

This indicates that as we haven't updated **hosts**, there is no need to process it. If you want to process all of the templates without checking their timestamps, then just give **ttree** the **-a** option.

The template itself is pretty simple: there's not much in it that we didn't cover last month. The only interesting thing is that we are using the *datafile* plugin twice – once to open the mail 'networks.txt' file and then again within the loop to open each individual network's data file. In this second usage, we have to use the more explicit syntax **\$(network.netname)** to reference the network's name. Without it the parser would have had little chance of realising what we meant. The Template Toolkit parser can usually work out what you want it to do, but in rare cases (like this one) it needs a bit of a hint.

If you look in the `./out` directory, you will see the results of processing our template which should look like this:

```
#
# /etc/hosts
#
# Network: sunnydale
# IP: 192.168.1/24
192.168.1.1 buffy slayer
192.168.1.2 willow witch
192.168.1.3 xander
192.168.1.4 spike
# End of network sunnydale
# Network: los_angeles
# IP: 192.168.2/24 i
192.168.2.1 angel
192.168.2.2 cordelia cheerleader
192.168.2.3 wesley
192.168.2.4 spike
# End of network los_angeles
# End of networks
```

Some more *ttree* options

Let's fill in a few more options in our `.ttreerc` file, so it looks like this following example:

```
verbose
recurse
ignore = \b(CVS|RCS)\b
ignore = ^#
copy = \.png$
copy = \.gif$
```

THE SUNNYDALE NETWORK

Two subnets and a single bridge

Throughout this tutorial series, we'll be using examples based on a very simple network. The network has two subnets and a single bridge between them.

The first subnet is 192.168.1/24. It has three computers called buffy (192.168.1.1), willow (192.168.1.2) and xander (192.168.1.3). The second subnet is 192.168.2/24. It also has three computers called angel (192.168.2.1), cordelia (192.168.2.2) and wesley (192.168.2.3). The bridge between the two networks is called spike and it has the two IP addresses 192.168.1.4 and 192.168.2.4. The illustration below shows how this network is connected.

Data about this network is held in a number of files. 'networks.dat' contains details of the subnets and 'net_sunnydale.txt' and 'net_los_angeles.txt' contain details of the computers on each of the subnets.

The contents of the files are:

```
# networks.txt
netname : number
```

```
sunnydale : 192.168.1/24
los_angeles : 192.168.2/24
```

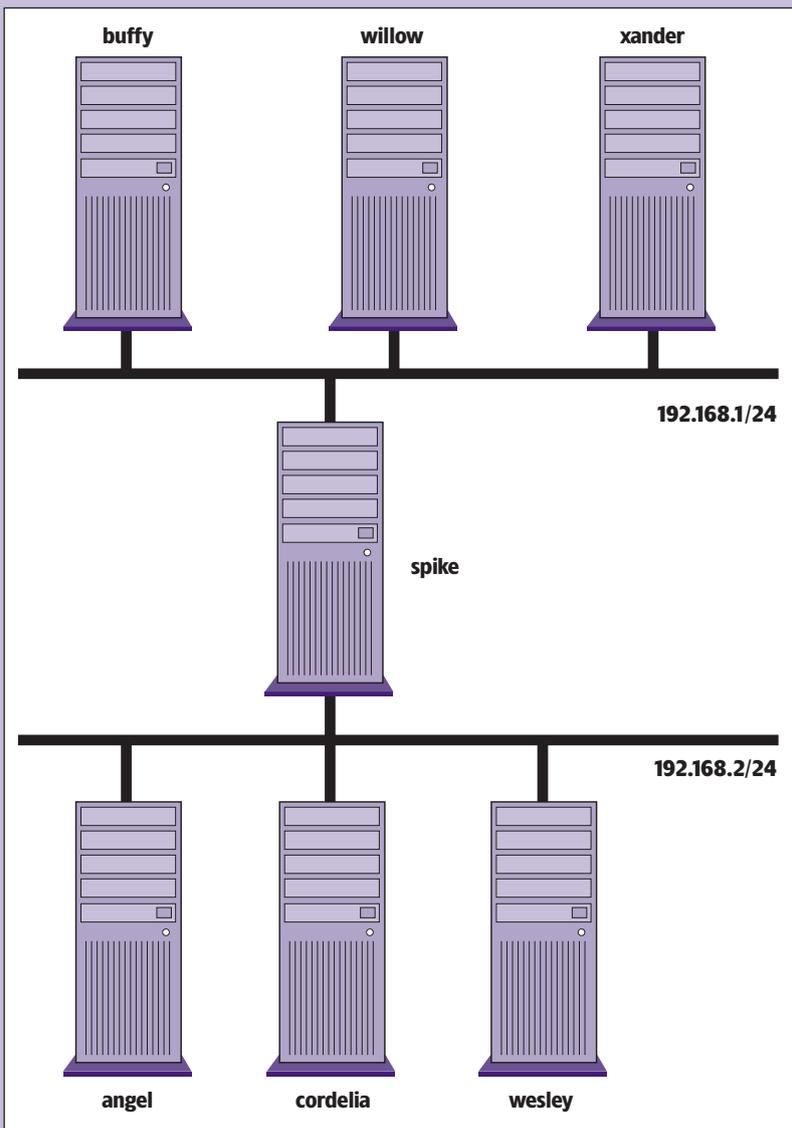
```
# net_sunnydale.txt i
```

```
IP : hostname : alias
192.168.1.1 : buffy : slayer
192.168.1.2 : willow : witch
192.168.1.3 : xander
192.168.1.4 : spike
```

```
# net_los_angeles.txt
```

```
IP : hostname : alias
192.168.2.1 : angel
192.168.2.2 : cordelia : cheerleader
192.168.2.3 : wesley
192.168.2.4 : spike
```

As we saw last month, these files are deliberately designed to be in the default format used by the Template Toolkit *datafile* plugin, but it would be equally simple to get the data from other file formats, XML documents or even a database.



```
◀◀ src = ./in
```

```
dest = ./out
```

```
lib = ./lib
```

The **recurse** option tells *ttree* to look in any subdirectories below your source directory and to recreate the same directory structure under the destination directory. The **ignore** option lists files that should never be processed. The arguments to this option are interpreted as Perl regular expressions and filenames that match the regular expression are ignored. In this example, the first **ignore** line matches CVS or RCS, thereby removing any sourcecode control files from consideration. The second line matches files that start with **#**, thereby ignoring *Emacs* backup files. If you're a *Vi* user, you might like to replace that with

```
ignore = ~$
```

The next option is **copy**. This lists files that are simply copied from source to destination without processing. Here we are copying png and gif files. The final option defines a library directory. This is an additional directory where *ttree* will look for templates. This is often used to store templates that are included in other templates and that aren't intended to be processed on their own.

Two other useful options that we won't be using in this example are **accept** and **suffix**. **accept** – as you might have guessed – is the opposite of **ignore**, as it defines the set of files that will be processed. You normally only need to use one of **ignore** or **accept**. **suffix** gives you a way to change the extension of files as they are being processed. For example, you might want to have a standard extension of **.tt** for templates, but convert that to **.txt** for the output files. In that case you could have a **suffix** option which looked like

```
suffix tt=txt
```

Creating More Files

The main advantage that *ttree* has over the simpler *tpage* is that it processes a complete directory of templates in one go. So far, our example only processes one template. So let's add another.

Another file that could potentially be derived from our network definition data is a *BIND* configuration file, so here is a template that could be used to create such a file.

```
[% PROCESS config;
file = 'db' _ main_domain;
FILTER redirect(file);
PROCESS soa domain=main_domain -%]

[% main_domain %] IN NS [% dns %].[% main_domain %]

; Hosts
[% USE networks = datafile('data/networks.txt');
FOREACH network = networks;
USE hosts = datafile("data/net_${network.netname}.txt");
FOREACH host = hosts -%]
[% host.hostname _ ' ' _ main_domain _ ' ' | format('%-32s') %]
IN A [% host.IP %]
[% IF host.alias -%]
[% host.alias _ ' ' _ main_domain _ ' ' | format('%-32s') %] IN
CNAME [% host.hostname %].[% main_domain %].
[% END;
END;
END -%]
STTL [% ttl %]
```

NEXT MONTH

So far we haven't actually had to write any Perl code at all in order to process our templates. Next month we'll look at how you use the Template Toolkit from within Perl programs.

SPECIFYING TREE OPTIONS

Controlling your processing runs

In this article we have mainly been controlling *ttree* by putting option definitions in the *.ttreerc* file. It's also possible to pass options on the command-line. This can be useful if you want to override a value from your *.ttreerc* for one or two processing runs and it's not worth the effort to edit *.ttreerc*. You can get a complete list of these options from **ttree -h** but here is a list of the values that we have used in this tutorial.

General options:

-a (--all)	Process all files, regardless of modification
-r (--recurse)	Recurse into sub-directories
-v (--verbose)	Verbose mode
-s DIR (--src=DIR)	Source directory
-d DIR (--dest=DIR)	Destination directory
-l DIR (--lib=DIR)	Library directory
-f FILE (--file=FILE)	Read named configuration file

File search specifications:

--ignore=REGEX	Ignore files matching REGEX
--copy=REGEX	Copy files matching REGEX
--accept=REGEX	Process only files matching REGEX

One thing that you'll notice immediately is that because this template uses a lot of directives we have started to combine multiple directives within one tag set. The Template Toolkit parser allows you to do this as long as you separate the directives with semicolons.

This template is far more complex than anything that we've seen before, so it's worth going through it in some detail. It starts by processing another template called **config**, which is shown below.

```
[% main_domain = 'whedon.example.com'
ttl = '3h'
dns = 'xander'
hostmaster = 'hostmaster.whedon.example.com'
-%]
```

This is a good example of a library template. All it does is define some variables that we will need elsewhere. We don't want to put it in the source directory as then it will be processed by *ttree* and we will end up with an extra unnecessary output file. Therefore you should put the config template in the *.lib* directory.

The next thing the template does is to create a new variable **file** which contains the name of the required output file. In this example, **file** will get the value **db.whedon.example.com**. We do this because we will eventually want to create a number of *BIND* configuration files and it will be nice to create them all using the same input template. To actually write the output to the correct file, we use the **redirect** filter. This takes one parameter which is a filename and writes the output from the filter to that file. Everything from the opening **FILTER** directive to the matching **END** will end up in the new file. In this example, the **END** that matches our **FILTER** right at the end of the template, so everything is written to the given file.

The next directive processes another external template called **soa**. This template provides the 'start of authority' block for the *BIND* file. Again, we've created a separate template as we would like to use it from several different templates. The template is shown below:

```
STTL [% ttl %]
```

```
[% domain %]. IN SOA [% dns %].[% domain %]. [%
hostmaster %]. (
  [% serial %] ; serial
  3h ; refresh
  1h ; retry
  1w ; expire
  1h ) ; caching TTL
```

This is another file which we don't want *tree* to process, so once again we put it in the `/lib` directory. Notice that the `soa` template uses a variable called `domain` and that this is passed in as a named parameter in the `PROCESS` directive.

The next directive in the template simply adds the NS record to the output file. It uses simple variable expansion that we've seen many times before.

Then we come to the part of the template which creates the **A** and **CNAME** records for the hosts on the network. This uses the same kind of logic that we used for the hosts file to loop through the data contained in the various data files and display the correct records. One nice touch is that we use the `format` filter to ensure that the domain name part of the record is always padded to the same length. Here we use the short syntax for the `FILTER` directive where the `FILTER` keyword is replaced by the pipe character (`|`). This makes it read a bit like a Unix filter command like `ls -l | sort`.

The template generates an **A** record for each host and a **CNAME** record for any aliases.

There's one thing missing from this description of the template. Sharp-eyed readers (and *BIND* experts) will have noticed that the `soa` template uses a variable called `serial` and that hasn't been defined anywhere. As the serial number needs to be incremented for each version of the configuration file, I thought that it was pointless to include it in any of the template files. You could, of course, include it in the `config` template, but you would need to remember to update it each time you processed the template. In my opinion, it's much easier to pass this value on the command-line to *tree*, and *tree* supports the same `--define var=value` option as *tpage* does. You can therefore process both of our templates with a command like this

```
$ tree --define serial=1
```

And you'll see that both the `hosts` template and the `db` template are processed.

Other *BIND* Files

Of course one `db` file doesn't make a complete *BIND* configuration. You'll need to define reverse lookup files for the `1168.192.in-addr.arpa` and `2.168.192.in-addr.arpa` domains as well as for the loopback domain. You'll also need the actual 'named.conf' file that pulls all of these files together. I don't have space to demonstrate creating all of these in this tutorial, but I hope I've given you some ideas on how you might go about it.

Template Complexity

This month's templates have been a lot more complex than the ones that we saw last month: it would be easy to argue that they were too complex! Part of this complexity comes from trying to do too much in a template. Templates should really only be concerned with presentation logic.

We can simplify our templates significantly if we use Perl to gather the data that we want to display, and only use the Template Toolkit language to control how we display that data. We'll look at how you do that in next month's tutorial. [LXF](#)



SUSE LINUX 9.1: The complete package for starters and professionals. Runs alongside Windows®!

SUSE LINUX
Personal 9.1
Now £ 24,95*

Finally I can switch
to LINUX easily
and completely.

SUSE LINUX Personal is the simple and secure entry into the operating system of the future. You'll get all the important applications to operate a complete workstation: Office, Multimedia, Mail, Internet and much more. It's time to change.

www.suse.com

 simply change

NETWORKING

The NET-SNMP package



PART 2 Following on from his series on IP networking, Dr Chris Brown continues his explanation to make sure you find that SNMP is indeed “simple.”

Last month, we looked in some detail at the theory behind the Simple Network Management Protocol (SNMP). Having first shown a description of the purpose of SNMP and the notion of agents and managers, we then looked at the MIBs, which define the structure of the information that an SNMP agent is able to return, and at OIDs (object IDs) which name variables within the SNMP object-tree in much the same way as pathnames identify files in the Linux filesystem.

If you made it to the end of that piece, you’ve probably had more than enough theory for the time being, and would relish something more practical to get your teeth into. We aim to please: so, this time we’re going to install and play with an Open Source set of SNMP tools called NET-SNMP, which you can find on the LXF coverdiscs of this very issue.

Naming the parts

As you will see from the program included on the discs, the NET-SNMP package consists of:

- 1 *snmpd*: an snmp agent for UNIX / linux
- 2 *snmptrapd*: a daemon to receive and log SNMP trap messages
- 3 A set of command line tools that are great for testing and can form the basis of an SNMP management system
- 4 A set of MIBS
- 5 Libraries to support writing your own management tools, or extending the *snmpd* agent
- 6 Manual pages for most of the above

```

Chris@dopey: m2011$ snmptranslate -Tp
oid(1)
---oid(1)
---oid(1)
  interface(1)
  ---oid(1)
  oid(1)
  ---oid(1)
    ipMIB(1)
    ---oid(1)
      sysName(1)
      sysLocation(1)
      sysContact(1)
      sysOID(1)
Chris@dopey: m2011$
  
```

Fig6 Output from `snmptranslate -Tp`.

A classic open-source installation

Before setting out to install the NET-SNMP package, you might want to check whether or not it’s already installed as part of your Linux distribution – it may well be. For example, try:

```
$ snmpget --version
```

to see if the command is found, and if so, which version it is. If you have version 5.0 or above, you can probably stick with that. If you have an earlier version, you should upgrade.

The current version, as described in this tutorial, is 5.1 and it’s available from net-snmp.sourceforge.net as a compressed tar archive. The file you want is `net-snmp-5.1.targz`. (Follow the ‘Download’ link from the home page if you’re not using the version on the coverdisc). After uncompressing and unpacking the archive, the build follows the standard `./configure && make && make install` sequence that will be familiar if you’ve previously built Open Source software on Linux. Like most configure scripts, this one takes a shedload of options that control where the various pieces of the package will be installed and which features will be enabled. The default configuration is probably fine for a first build of the package, and will install everything under `/usr/local`. After the usual screenfuls of “checking for..” messages, the configure script will ask a few questions, something like this:

1. Default version of SNMP to use (3): 2
2. System Contact Information (chris@): chris@dopey
3. System Location (Unknown): Back Office
4. Location to write logfile (/var/log/snmpd.log):
5. Location to write persistent information (/var/net-snmp):

This dialog is interspersed with lots of useful prompts which we’ve omitted here due to space limitations. If you just want to play with SNMP and aren’t concerned about high security, I’d suggest having the tools default to version 2 of the protocol (see line 1 above). In any case, this is easy to over-ride on the command-line when you start up the individual tools. At lines 2 and 3, just supply appropriate values. These entries provide default values for the `sysContact` and `sysLocation` variables returned by your *snmpd* agent and can be over-ridden later in the agent’s config file. Lines 4 and 5 are fairly self-explanatory, and, having no reason to do otherwise, I just accepted the defaults.

After completing the configure script, run `make` followed by `make install`. This last command must be run as root because it copies the files into their correct system directories.

Once installation is complete, we can take a look around.

Table 1 below shows some of the key directories, and what's installed there:

TABLE 1: KEY DIRECTORIES

Directory	Contents
<code>/usr/local/sbin</code>	The SNMP agent <code>snmpd</code> and <code>snmptrapd</code>
<code>/usr/local/man</code>	The manual pages
<code>/usr/local/bin</code>	The command-line tools
<code>/usr/local/share/snmp/mibs</code>	The MIB definitions

Configuring the snmpd agent

Whenever I install new software, I always want to demonstrate basic functionality as soon as possible. After that I relax, and I'm willing to explore more thoroughly and slowly. I see this as the geek equivalent of a child ripping the wrapping paper off a gift. (My wife just mutters something about the delights of "delayed gratification" – whatever THAT is...). Anyway, if you have any SNMP-enabled devices on your network, you can go ahead right now and try out some of the command-line tools like `snmpget` to query them. But I'm going to assume that you haven't, so we'll need to get `snmpd` running in order to have an agent that we can query. For `snmpd` to function, it needs a configuration file, `/usr/local/share/snmp/snmpd.conf`. The distribution includes a perl script called `snmpconf` which will create an `snmpd.conf` file by asking a series of questions. In modern parlance, I suppose it would be called a wizard, though to be honest it's almost harder to understand and answer the wizard's questions than it is to hand-edit the config file, and as wizards go I see it as more of a Ron Weasley than a Gandalf.

In fact, `snmpd` will run with a config file containing just the following two lines:

```
rocommunity rocom
rwcommunity rwcom
```

Here, **rocommunity** and **rwcommunity** are keywords used in the file, whereas **rocom** and **rwcom** are just names I made up. These two lines define two 'communities' that can access the agent, one in read-only mode, and one in read-write mode. If you remember back to last month's tutorial, a community is essentially just a password known to the agent that allows access to the MIB.

With these two lines in the config file, it's time to start up the agent. As root, run the command

```
# /usr/local/sbin/snmpd
```

There's no need for an **&** at the end of the line, as the agent will background itself automatically. Now run

```
netstat -au
```

and check for a line that has ***:snmp** in the Local Address column. This tells us that we have a UDP socket created on `snmp`'s standard port number (161). If you don't, you'll need to investigate why `snmpd` won't start up.

Assuming that all's well, we can try running a simple query on our agent, like this:

```
$ snmpget -v 1 -c rocom localhost 1.3.6.1.2.1.4.0
SNMPv2-MIB::sysContact.0 = STRING: chris@dopey
```

Here, **1.3.6.1.2.1.4.0** is the OID of the `sysContact` variable, and you should receive back the string you entered to define the system contact information when you ran the configure script.

At this point, we've demonstrated basic operation of the agent and the tools, and can take time to look around more fully.

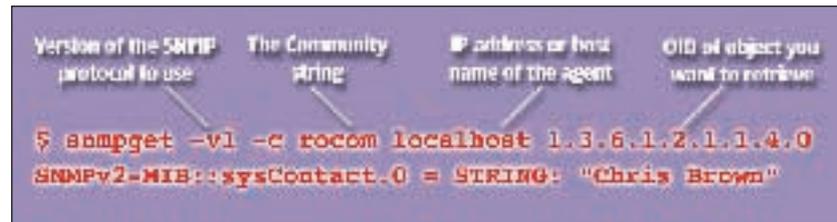


Fig7 Dissecting an `snmpget` command.

Exploring the MIBs

A good place to start exploring might be to look at the MIB definitions that are included in the NET-SNMP package. The MIBs are in text files in the directory `/usr/local/share/snmp/mibs`, and have names like `IF-MIB.txt`. I suggest that you start by looking at the three MIBs `SNMPv2-MIB`, `IF-MIB` and `IP-MIB`. These files are written in the language ASN.1 and we talked about them briefly last month. We saw that each variable defined in the MIB has a name, a data type, an access level, and a text description.

While it's worth seeing how the MIBs are written, reading large pieces of them directly is about as exciting as studying the telephone directory, or, if I'm allowed a moment's heresy, Appendix F of *Lord Of The Rings*. Also, it's quite tedious to read the MIBs and manually reverse engineer from them the object tree which they define. There are tools in the NET-SNMP package that make the job easier. In particular, the `snmptranslate` program reads in all the MIBs, locates a specified variable, and prints it out in a variety of formats. For further understanding, let's take a look at `snmptranslate` in action.

A useful option for `snmptranslate` is **-Td**, which prints all known information about an object in a format resembling the original MIB definition. Some example code follows, after a brief explanation. Here, we're querying `snmptranslate` for the object ID `1.3.6.1.2.1.4.2` which happens to be the default time-to-live field for the IP layer .. the same variable we used as an example in part 1 of this tutorial. In fact, if you compare the entry in the MIB that we showed last month with the output from `snmptranslate` below, you'll see they're very similar:

```
$ snmptranslate -Td 1.3.6.1.2.1.4.2
IP-MIB::ipDefaultTTL
ipDefaultTTL OBJECT-TYPE
-- FROM IP-MIB, RFC1213-MIB
SYNTAX INTEGER (1..255)
MAX-ACCESS read-write
STATUS current
DESCRIPTION "The default value inserted into the Time-To-Live field of the IP header of datagrams originated at this entity, whenever a TTL value is not supplied by the transport layer protocol."
 ::= { iso(1) org(3) dod(6) internet(1) mgmt(2) mib-2(1) ip(4) 2 }
```

One feature of the `snmptranslate` command-line that's common to most of the NET-SNMP tools is the range of formats available for specifying an OID. The format shown above, with a full numeric OID, is one possibility. You can also use the full textual form of the name, like this:

```
$ snmptranslate -Td iso.org.dod.internet.mgmt.mib-2.ip.ipDefaultTTL
```

which is less cryptic but tediously long. You can also specify a MIB name and a variable name, like this:

```
$ snmptranslate -Td IP-MIB::ipDefaultTTL
```



TUTORIAL Networking

Fig8 Walking the 'interfaces' table.

```

$ snmptranslate -Td -IR ipDefaultTTL
...
ifIndex: 1
ifDescr: eth0
ifType: Ethernet-CRM
ifMtu: 1500
ifPhysAddress: 00:0C:29:00:00:00
ifName: eth0
ifOperStatus: up
ifAdminStatus: up
ifLastChange: 1000000000
ifInDiscards: 0
ifInErrors: 0
ifInPkts: 1000000000
ifInSpeed: 1000000000
ifInUcastPkts: 1000000000
ifInMulticastPkts: 0
ifInBroadcastPkts: 0
ifOutDiscards: 0
ifOutErrors: 0
ifOutPkts: 1000000000
ifOutSpeed: 1000000000
ifOutMulticastPkts: 0
ifOutBroadcastPkts: 0
...

```

Finally, there's an **-IR** option which enables 'random access' into the MIBs and enables you to specify a simple, unqualified name, like this:

```
$ snmptranslate -Td -IR ipDefaultTTL
```

Of course, it's possible that the name may appear in more than one MIB. The manual pages are a bit vague about what happens in this situation.

Another useful option for **snmptranslate** is **-Tp**, which prints a specified portion of the object tree using clunky but effective character graphics. **Fig6** on the previous page gives you an idea what this looks like; it's **snmptranslate**'s rendition of the object tree (or at least, the first 20 lines of it) starting from the OID **.1** – ie the very top of the tree. The full output from this command is quite extensive (about 1800 lines) and of course this is only from the MIBs that ship as part of NET-SNMP. If you have vendor-specific MIBs installed, the output might be much longer.

Before we turn our attention away from **snmptranslate**, let's just emphasise that it's purely a tool for interpreting the MIBs. It does not query an SNMP agent.

Talking to the agent

The simplest command-line tool for querying an agent is **snmpget**. **Fig7** on the previous page dissects a simple **snmpget** command. The community string defined via the **-c** option is passed as part of the request and used by the agent to authenticate the request. **snmpget** supports the same formats for specifying the OID as we saw for the **snmptranslate** command. In addition, there are options that control the output format – see the **snmpcmd** manual page for details of these. In all our examples this month we've used **localhost** as the location of the agent, but you can specify an IP address or the name of any host or device where there's an SNMP agent running.

The **snmpset** command lets you change the value of those (relatively few) **snmp** variables which have read/write access. In the dialog below we first query the **snmpd** agent for the value of the **sysContact** variable. Then we run **snmpset** to change the value, and finally we run **snmpget** again to verify the change. Notice that for the **snmpset** command we specified the community string of a community (**rwcom**) with read/write access to the agent.

```
$ snmpget -v 2c -c rocom localhost SNMPv2-
MIB::sysContact.0
```

```
SNMPv2-MIB::sysContact.0 = STRING: chris@dopey
```

```
$ snmpset -v 2c -c rwcom localhost SNMPv2-
MIB::sysContact.0 = fred
```

```
SNMPv2-MIB::sysContact.0 = STRING: fred
```

```
$ snmpget -v 2c -c rocom localhost SNMPv2-
MIB::sysContact.0
```

```
SNMPv2-MIB::sysContact.0 = STRING: fred
```

Next we come to **snmpgetnext** – not a particularly useful command by itself, being essentially just a wrapper around the **GET-NEXT** operation. It works much the same as **snmpget** except that it returns the object following the one you asked for. We spent a long time understanding the notion of **GET-NEXT**, and lexicographic ordering of the object tree, and the idea of 'walking' the MIB in the last tutorial. (Ha! ... Now I bet you wish you'd got there before the newsagent sold out – backissues are available on page 95). We saw that repeated use of **GET-NEXT** allows you to easily traverse sections of the MIB, especially tables, without knowing in advance exactly what variables you'll find.

A much more useful command that makes repeated use of **GET-NEXT** is **snmpwalk**, which prints out all the entries below a given node in the subtree, in lexicographic order. **Fig8** shows an example of walking the MIB starting at the 'interfaces' node, which you'll find near the bottom of **Fig2** and in more detail in **Fig5**, both illustrations in last month's tutorial. We haven't shown the whole thing, but we've shown up to the **ifPhysAddress** column of the table.

There are a couple of things to notice here. First, we're traversing the entire table, working down the first column (**ifIndex**) then the second (**ifDescr**) then the third (**ifType**) and so on. The table has four rows, suggesting that the machine I ran it on has four network interfaces. In fact, the first row is the loopback interface, the second is the real network card, and the third and fourth are virtual interfaces created by the **VMware** software I run on this machine.

Using **snmpwalk** to traverse tables sure is a big improvement on retrieving each item separately with **snmpget**, but it displays tables in a funny order and a not particularly helpful format. In fact, there's another command, **snmptable**, which re-orders the output and lays it out better. **Fig9** shows an example, showing essentially the same information as the **snmpwalk** command above. The **-Cw 80** option sets a column width of 80 characters and forces **snmptable** to break the output into several parts, of which only two are shown here.

There are a few more specialised commands for querying an **snmp** agent, including **snmpdf**, which queries for disk space usage on the agent and produces output similar to the **df** utility, and **snmpnetstat**, which queries network connections on the target machine and produces output similar to **netstat**. However, it's probably time to turn our attention back to the **snmpd** agent.

```

$ snmptable -Cw 80 localhost SNMPv2-MIB::ifTable
ifIndex  ifDescr          ifType  ifMtu  ifPhysAddress  ifName  ifOperStatus  ifAdminStatus  ifLastChange  ifInDiscards  ifInErrors  ifInPkts  ifInSpeed  ifInMulticastPkts  ifInBroadcastPkts  ifOutDiscards  ifOutErrors  ifOutPkts  ifOutSpeed  ifOutMulticastPkts  ifOutBroadcastPkts
1         lo                 6        1500   00:00:00:00:00:00  lo      up             up              1000000000    0           0           1000000000  1000000000  0           0           0           0           1000000000  0           0           0
2         eth0               6        1500   00:0C:29:00:00:00  eth0    up             up              1000000000    0           0           1000000000  1000000000  0           0           0           0           1000000000  0           0           0
3         veth0              6        1500   00:00:00:00:00:00  veth0   up             up              1000000000    0           0           1000000000  1000000000  0           0           0           0           1000000000  0           0           0
4         veth1              6        1500   00:00:00:00:00:00  veth1   up             up              1000000000    0           0           1000000000  1000000000  0           0           0           0           1000000000  0           0           0

```

Fig9 Output from **snmptable**.

RESOURCES

If you would like more information about NET-SNMP and SNMP in general, take a look at the following websites:

<http://net-snmp.sourceforge.net>

Home site for the NET-SNMP package

www.simple-times.org

Hosts a rather irregularly published newsletter about SNMP. There's some interesting material here but the most recent newsletter is dated December 2002.

www.mibdepot.com

Provides a large, downloadable database of MIBS and a useful search engine

www.simpleweb.org

Lots of good stuff here – links to standards organisations, commercial and free SNMP software, tutorials, all the relevant RFCs, and a great browser-based interface that we featured last month that allows you to explore the MIBs.

Meanwhile, Back at the Agent...

With the minimal two-line config file we created earlier, **snmpd** will answer queries for objects in the **system**, **ip**, and **interfaces** groups; and in the **tcp**, **udp** and **icmp** groups, among several others. However, there are other directives you can include in **snmpd.conf** which extend the agent and allow it to monitor various aspects of the health of your Linux system.

For example, the **proc** directive makes **snmpd** monitor the number of running instances of a given process (say, **httpd**), and report if the number falls outside specified minimum and maximum bounds. The **procfix** directive lets you specify a command to be run to take remedial action if there's a problem. The command is triggered by writing to a specific SNMP variable, allowing SNMP management stations to be pro-active in managing the health of network servers. Likewise the **disk** directive makes **snmpd** monitor disk usage of specified disk partitions and the **load** directive makes it monitor the load average. Run

```
man snmpd.conf
```

for more details of these directives.

Extending the Agent

One of the features of **snmpd** that makes it appeal to users is that it's extensible – either by writing your own modules in C, or by defining external programs to be run which can contribute data to the MIB. Writing our own MIBs and modules is beyond the scope of this tutorial, but extending the agent with external programs is well within our grasp.

It works like this. The agent maintains a table at an OID within the **private enterprise** part of the object tree. It's called **extTable** and is defined in the MIB UCD-SNMP-MIB. It's rooted at 1.3.6.1.4.1.2021.8. To configure the agent to build the rows of the table you put **exec** directives into the **snmpd.conf** file which name external commands to be run. (Each **exec** directive contributes one row to the table.) The agent captures the exit status and the first line of the standard output of the commands, and puts them into the table. For this example, we'll extend the agent to report the version of the kernel it's running on (as reported by **uname -r**) and the total number of **httpd** processes running (as reported by **ps ax | grep httpd | wc -l**). I added two lines into **snmpd.conf** like this:

```
exec kernel_version /bin/uname -r
```

```
exec httpd_count /usr/local/bin/hcount
```

The first line causes the output from **uname -r** to appear in the first row of the table and the second causes the output from a one-line script called **hcount** to appear in the second row of the table. Notice that I've specified full pathnames for these commands. The script simply contains:

```
#!/bin/bash
```

```
ps -ax | grep httpd | wc -l
```

Why do I need to use a script instead of just putting the **ps | grep | wc** pipeline directly after the **exec** directive? Well, it turns out that **snmpd** executes these commands directly, so they are never seen by a shell. Consequently, things like pipes and filename wildcards won't be recognised. Notice that this is an 'expensive' solution, as it requires four extra processes to be started up (one to run the script and one each for **ps**, **grep** and **wc**) every time the table is queried.

With these two directives and the **hcount** script in place, I can stop and restart the **snmpd** agent and use **snmpwalk** to print the two-row table that the agent has created.

extIndex (1)	extName (2)	extCommand (3)	extResult (100)	extOutput (101)
1	kernel_version	/bin/uname -r	0	2.4.18-14
2	httpd_count	/usr/local/bin/hcount	0	10

Fig 10 **extTable**.

```
$ snmpwalk -v 2c -c rocom localhost 1.3.6.1.4.1.2021.8
```

```
UCD-SNMP-MIB::extIndex.1 = INTEGER: 1
```

```
UCD-SNMP-MIB::extIndex.2 = INTEGER: 2
```

```
UCD-SNMP-MIB::extNames.1 = STRING: kernel_version
```

```
UCD-SNMP-MIB::extNames.2 = STRING: httpd_count
```

```
UCD-SNMP-MIB::extCommand.1 = STRING: /bin/uname -r
```

```
UCD-SNMP-MIB::extCommand.2 = STRING:
```

```
/usr/local/bin/hcount
```

```
UCD-SNMP-MIB::extResult.1 = INTEGER: 0
```

```
UCD-SNMP-MIB::extResult.2 = INTEGER: 0
```

```
UCD-SNMP-MIB::extOutput.1 = STRING: 2.4.18-14
```

```
UCD-SNMP-MIB::extOutput.2 = STRING: 9
```

This table is also illustrated in **Fig 10**. Columns **2** and **3** of the table simply show the arguments to the **exec** directive. The last two lines of output from **snmpwalk** (corresponding to the right-hand column of the table) are the most interesting as they show the standard output from the commands, telling us (in this example) the version of the Linux kernel we're running, and the number of **httpd** processes.

The NET-SNMP libraries

NET-SNMP includes a set of libraries to support writing your own management code (in C), and to support the writing of modules that extend the **snmpd** agent and allow you to define and implement your own MIBs.

There is some tutorial information and example code on the www.net-snmp.org website, and some of the library functions even have man pages. There's a MIB compiler called **mib2c** that will generate template agent code from your own MIB definitions. The documentation in this area is not very complete, though, and is extremely terse, so if you're thinking of writing your own manager or agent code, be prepared for a fairly rough ride and some digging around in the header files and source code. Nonetheless, if you're looking to implement an SNMP agent on some Linux-based network appliance you've developed, it sure beats starting from scratch!

Summary

The NET-SNMP package provides a working, extensible SNMP agent for Linux-based systems along with a useful set of command-line tools. These tools are fairly basic and don't aspire to provide the comprehensive reporting and fancy graphical displays you'll find in commercial SNMP management products. Still, they make an excellent starting point for exploring SNMP, understanding the protocol, and testing other SNMP implementations. If you've any comments or suggestions about this tutorial series, please send them in to the addresses listed on page 12. [LXF](mailto:lxf@linuxformat.co.uk)

NEXT MONTH

Having got our hands dirty by installing the NET-SNMP package, next issue's tutorial will be throwing you in at the deep end, now you've had a chance to get your head around the basics...

Answers

If you are really stuck and the HOWTOs yield no good result, why not write in? Our resident experts will answer even your most complicated problems!

Our experts

Whatever your question is, we can find an expert to answer it – from installation and modem woes to network administrations, we can find the answer for you – just fire off a letter or email and it'll all be taken care of.

LXF answers guy

David Coulson is a networking and security guru with plenty of sysadmin experience to boot.



Nick Veitch is the editor of the magazine, and answers your easy questions! Or indeed anything to do with *Grub*, *LILO*, *netatalk*, vi...



Hans Huberland is Rackspace Managed Hosting's Linux expert. Send any Linux system admin questions to sysadminqa@rackspace.co.uk



Failure in X

Q I recently suffered a failure of the X server. It wouldn't start (having previously worked fine for many months) and gave an error message saying that it was unable to init the default font. This may have been triggered by the / partition becoming full as determined by using *df*. I have managed to get X running again by editing *XF86Config* and adding specific paths to some of my fonts but originally the only path was `unix:/7100`. and this seems to be where the problem lies. If I try to run *xfs* from the command line I get an error message that says:

```
_FontTransmkdir: Owner of
/tmp/.font-unix should be set to root
FontTransSocketUNIXCreate
Listener: ...SocketCreateListener()
failed
_FontTransMakeAllCOTSServer
Listeners: server already running
xfs error: Cannot establish any
listening sockets
Aborted
```

The `/tmp/.font-unix` directory is in fact shown as being owned by *xfs*. I have reloaded *xfs* from the RPM file on the source CD to no avail. The *xfs* config file has a catalogue section that lists the locations of all my fonts, so why is *xfs* failing? Can you give me any pointers as to where I should look next?

Also, how can I print MAN and other shell pages? Using either the `>` or `!` redirection to LPR doesn't work and there is no print option in either the shell or *man* options. My printer is operated via CUPS

Colin Ager, Garboldisham, Norfolk

A It's quite possible that you corrupted your filesystem when it became full, creating some issues with programs which you

run. The first step should be to boot your Linux system into single-user mode, then manually run *fsck* across the filesystem which was full. Once this has been done, you may wish to remove then reinstall the *xfs* package in order to ensure the binaries and init scripts are correct.

Filling an ext2 or ext3 filesystem often creates a lot of problems, as blocks can become overwritten and files can end up being mashed together. Once *fsck* is run, it does resolve many of these problems, however it's not improbable that damage has been done to important files. As we don't know exactly how your drives are partitioned, it's possible that binaries and library files are corrupted in addition to any scripts existing in `/etc`.

To print a man page, we can simply pipe it through the *groff* interpreter that is used by the *man* utility to convert the document into something manageable. If we know the full path to the manual page, you can do:

```
groff -Tps -man /usr/share/man/
man1/cp.1 | lpr
```

Of course, one will need to pass the appropriate

```
-Pprinter
```

switch to *lpr* should you not be using the standard *lp* printer.

Mandrake mishaps

Q I have enjoyed reading your magazine for the last two years. However, I am at present having some problems with my Mandrake 9.2 installation using the Linux Format CDs. I installed the software as an upgrade to Mandrake 9 (the public release version). The system boots fine and runs fine (needs some tweaking of the sound subsystem). However, I cannot update software using *Mandrake Update* or install software from the CDs cleanly with the

RPMDrake application. Lastly, I have a lot of clutter from update and installation choices because it appears to be referencing items from MDK9 as well as MDK9.2.

1 *Mandrake Update* returns the following error when trying to upgrade to *coreutils-5.0-6.1.92mdk.i56*:

```
6.1.92mdk.i56:
```

```
Some packages can't be installed.
Sorry, the following package(s) can't
be selected:
coreutils-5.0-6.1.92mdk.i56 (due to
unsatisfied rpmlib(PayloadFiles
HavePrefix) <= 4.0-1)"
```

2 Using *RPMDrake* to install *proftpd*, The following packages have bad signatures:

```
proftpd-1.2.8-4mdk.i586.rpm:
Invalid signature ((SHA1) DSA sha1
md5 (GPG) (MISSING KEY)
GPG#70771ff3 NOT OK)
```

I am then asked: "Do you want to continue installation?" Of course I do! Any advice you have on how to clean up the installation would be appreciated. Thanks!

Mike Caron, USA

A Your first error would suggest that you do not have the most up to date version of *rpmlib*, which *coreutils* is a dependency of. You can verify the currently available version of *rpmlib* on Mandrake's FTP server and manually install it. Likewise with the GPG problem, the key that you need should be distributed as part of the current packages.

At a first glance, it would appear that your RPM packages have not been updated correctly. You may wish to start by installing them from the Mandrake 9.2 CD and seeing if that solves any of your problem. Of course, upgrading the RPM packages using *rpm* may present a few issues, so if there are any errors, it's best to avoid forcing the packages in any way.

Welcome to Linux

Q Hi, I am a beginner – or at least I would be, if I could begin! I am not new to PCs, having worked with computers since the days of Commodore 64 *et al.* My present PC is based on the ASUS A7V8X motherboard with a AMD Athlon XP 2200+ running at 1.8GHz. It has 1GB of RAM DDR 333/PC2700 and the operating system is Windows XP Home.

I first tried loading Linux from the disks supplied with the *Linux Desktop Starter Kit* by John Lathrop without success. Thinking the limitation might be due to the fact that this is a 'Publishers Edition' of Linux, I bought *Red Hat Linux 9 Bible* – which included three disks and purported to be 'the full works'. In both cases, I just could not load the software. The approach is simple: boot Windows, change BIOS to boot from CD, put Installation CD in the drive, switch off, then reboot, and the system



Red Hat no longer maintains a free distro, Fedora is a fork of the RH9 tree and is supported by Red Hat developers.

A QUICK REFERENCE TO: Asterisk

With the advent of much faster bandwidth services to homes and businesses, delivering more than simply Internet access is now a realistic option for many. While we're not quite there for high-bandwidth video services, voice is a realistic option for delivering over IP (VoIP). Avoiding the local telephone company allows us to route our voice to different VoIP providers, making for more cost effective branch offices, long distance and international calling.

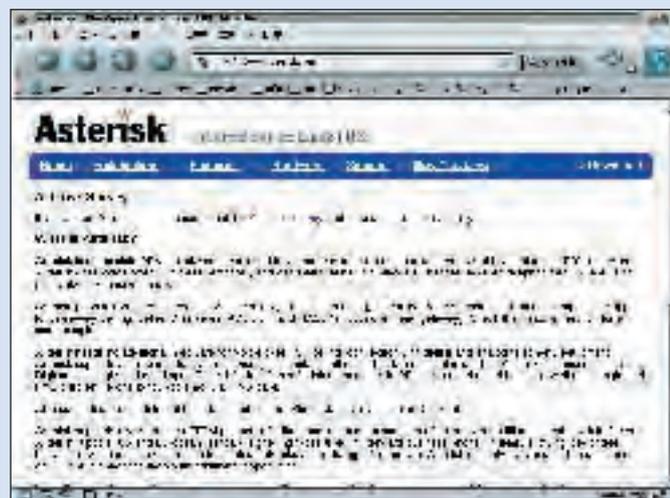
While there are commercial VoIP implementations from the likes of Cisco, there is an Open Source offering which provides all manner of voice gateway and PBX functionality for significantly less cost. *Asterisk* allows you to replace an existing PBX system, connected to POTS, or ISDN circuits, with a standard Intel box. IP services can be provisioned using one of the many voice phones available from manufacturers including Polycom, Cisco, Grandstream and Snom, depending upon costs and functionality needs.

Asterisk can trunk outgoing and incoming calls via VoIP providers, or using local circuits. Digium – a company that supports the Asterisk project – sells all the hardware required to interface with POTS, ISDN PRI and T1/E1. Connecting branch offices together can be done over IP, avoiding the need for expensive voice circuits when DSL or cable IP services are available.

Asterisk provides full PBX functionality, including call transfers, music on hold, conference calling, voicemail, call parking and almost anything else you can think of. This all makes it very easy for roaming users to access all of their voice services remotely and have calls forwarded in an easy to manage fashion.

We can have an extension call a variety of end points, including SIP phones, standard phones via a VoIP provider, an extension hanging off a channel bank, or a local number out of an ISDN circuit:

```
exten => 100,1,Dial(SIP/dcoulson&
IAX2/nufone/12165551234,10,tr)
```



Asterisk is a complete Open Source PBX implementation, offering VoIP and telco termination of calls.

```
exten => 100,5,Dial(IAX2/nufone/
12165551000,5,tr)
exten => 100,6,VoiceMail(u100)
exten => 100,7,Hangup()
exten => 100,102,VoiceMail(b100)
exten => 100,103,Hangup()
exten => 100,106,VoiceMail(b100)
exten => 100,107,Hangup()
```

In this example, we can make *Asterisk* call our SIP phone and our home number at the same time for ten seconds. Should this

fail, it will call our cell phone four five seconds before resorting to our voicemail. If when trying our SIP phone it receives a busy signal, then it jumps to voicemail immediately but with a different message saying we're busy.

Asterisk can be found at www.asterisk.org/ and information on VoIP providers, phone handsets and all the associated hardware is found in the very useful www.voip-info.org/ site.

FREQUENTLY ASKED QUESTIONS IP SECURITY

FAQ JUST HOW SECURE IS THE INFORMATION I TRANSPORT OVER IP?

Most protocols that use IP as a transport will transfer information in a plain text format, such as SMTP, POP and more complex services including LDAP and database sessions. Anyone creative enough to intercept traffic over these protocols could obtain passwords and other info used for authentication, as well as the actual data which has been transferred during the session.

FAQ DOES SSL SOLVE THIS PROBLEM?

SSL allows a specific session to be encrypted, assuming that the client and server on each end supports such a transport method. If the remote server does not support encryption of the desired protocol, then there is no way to force it to use SSL. It's common for SSL to be used for web services, but it's rare to find an ISP which offers mail services that utilise SSL.

If we are running our own mail server, it's fairly easy to implement SSL

services using a tool such as *sslwrap* or *stunnel*. Each of these allows us to wrap an existing plain text protocol within an encrypted tunnel. The server running does not need to know anything about the SSL connection, however the client does need support for SSL connectivity.

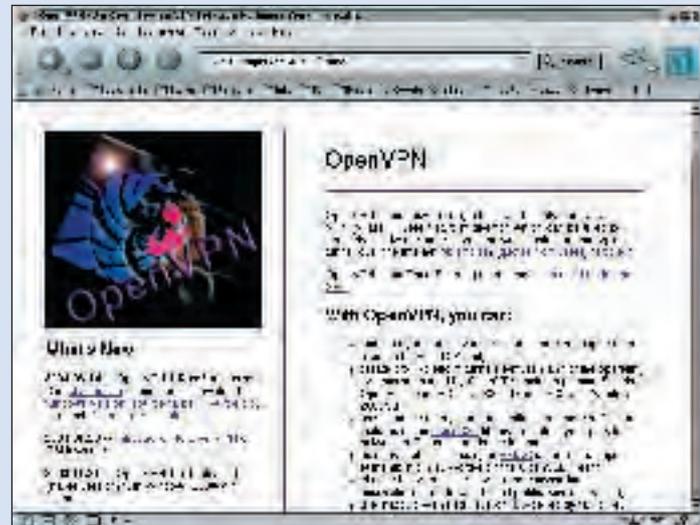
We can add entries to *inetd.conf* which provide SSL enabled services:

```
https stream tcp nowait
sslwrap.sslwrap /usr/sbin/tcpd
/usr/sbin/sslwrap -cert
/etc/sslwrap/server.pem -addr
127.0.0.1 -port 80
```

```
ssmtp stream tcp nowait
sslwrap.sslwrap /usr/sbin/tcpd
/usr/sbin/sslwrap -cert
/etc/sslwrap/server.pem -addr
127.0.0.1 -port 25
```

```
nntp stream tcp nowait
sslwrap.sslwrap /usr/sbin/tcpd
/usr/sbin/sslwrap -cert
/etc/sslwrap/server.pem -addr
127.0.0.1 -port 119
```

```
telnet stream tcp nowait
sslwrap.sslwrap /usr/sbin/tcpd
/usr/sbin/sslwrap -cert
/etc/sslwrap/server.pem -addr
```



OpenVPN allows us to create a secure private tunnel between two hosts, reducing the need for security support in the client and server.

```
127.0.0.1 -port 23
imap stream tcp nowait
sslwrap.sslwrap /usr/sbin/tcpd
/usr/sbin/sslwrap -cert
/etc/sslwrap/server.pem -addr
127.0.0.1 -port 143
```

```
ircs stream tcp nowait
sslwrap.sslwrap /usr/sbin/tcpd
/usr/sbin/sslwrap -cert
/etc/sslwrap/server.pem -addr
127.0.0.1 -port 194
```

```
pop3s stream tcp nowait
sslwrap.sslwrap /usr/sbin/tcpd
/usr/sbin/sslwrap -cert
/etc/sslwrap/server.pem -addr
127.0.0.1 -port 110
```

```
ftps-data stream tcp nowait
sslwrap.sslwrap /usr/sbin/tcpd
/usr/sbin/sslwrap -cert
/etc/sslwrap/server.pem -addr
127.0.0.1 -port 20
```

```
ftps stream tcp nowait
sslwrap.sslwrap /usr/sbin/tcpd
/usr/sbin/sslwrap -cert
/etc/sslwrap/server.pem -addr
127.0.0.1 -port 21
```

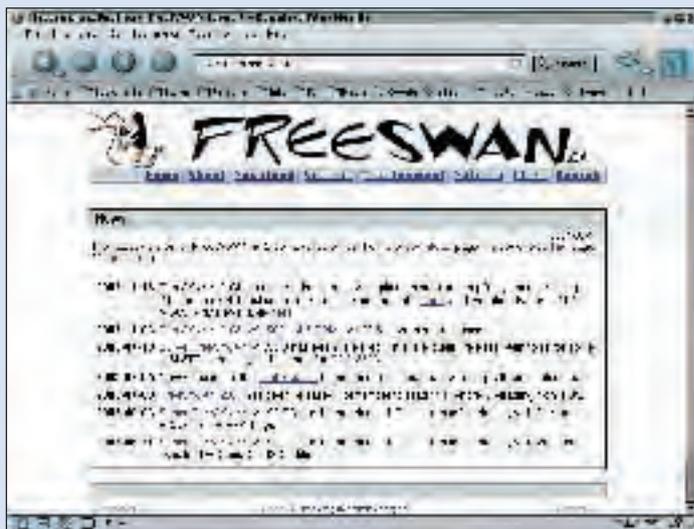
```
ldaps stream tcp nowait
sslwrap.sslwrap /usr/sbin/tcpd
```

```
/usr/sbin/sslwrap -cert
/etc/sslwrap/server.pem -addr
127.0.0.1 -port 389
```

FAQ HOW I CAN SECURE CONNECTIONS BETWEEN TWO HOSTS WITHOUT CLIENT OR SERVER SUPPORT.

To avoid the client/server support for encryption, we need to secure either the TCP or the IP layer of the connection. The *stunnel* utility allows us to run a local client to connect to a remote TCP port using SSL, then we can point the client which does not support SSL to a local port.

SSH also allows us to forward a TCP port over our secure SSH connection. We can define a remote port and host, along with a local port, and when we connect to the local port it will forward the connection across SSH to the remote host and port. Between our local system and the box we're SSHed to, the connection is secure. Out from the remote end, assuming the remote host is on a secured internal network, then the



FreeSWAN is a popular IPsec implementation for both 2.4 and 2.6 kernels, and offers secure connections to a variety of network devices.

starts from the CD and commences the Install process.

I respond to all the prompts and accept the default package for install. All looks OK until installation starts, after installing a number of files (anywhere between 14 and 86) the screen shows a message that

the installation has unexpectedly terminated and shuts down.

I tried various other types of Linux, including Mandrake and your recent pair of CDs for Fedora (LXF49, Jan 2004). After several attempts with Fedora (more than 20), trying variations on the response to the

prompts and reducing the number of packages to install, I eventually succeeded in installing a 'minimal' installation set. So I now have a basic working command-line system.

I thought it would be easy to upgrade by booting from the installation disk, selecting additional

packages and doing an update type installation. However, again I encounter the "unexpected termination" and shut down.

I have been wrestling with this since before Christmas, and have tried the Linux forums with little success and certainly no solution.

entire end to end connection will be secure. Everything that goes over the Internet will be transported over *SSH*, so no one malicious can sniff the data from it or inject new information into our session.

FAQ CAN I ENCRYPT ALL DATA BETWEEN TWO HOSTS?

Encryption at the IP layer is a little more complex, as it requires the use of an IP tunnel. There are a number of different IP tunneling tools available for Linux including *vtun*, *openvpn* and *FreeS/WAN*.

Both *vtun* and *openvpn* deliver IP information over an encrypted UDP or TCP tunnel, and require a significant amount of routing knowledge as one has to route a non-routeable network block over the tunnel. This is a great way to connect two internal and non-routeable networks together, or to create a virtual network between two hosts, or one host and an internal network.

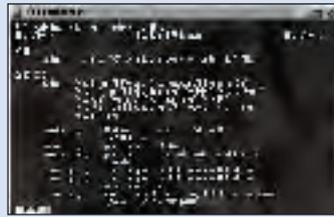
We can also use the ever-popular IPsec protocol. This is an encryption method that allows standard IP data to be transported between two hosts without any virtual network being created or the use of internal routes being used.

FAQ HOW DO I USE IPSEC WITH LINUX?

There are currently two implementations of IPsec for Linux. The most well known is *FreeS/WAN*, which is available for both Linux 2.4 and 2.6. The other is the version of IPsec which is currently included as part of the standard 2.6 kernel.

We're going to look at *FreeS/WAN*, as this is what most Linux people who want IPsec support will use. Information on using the 2.6 implementation can be found at <http://lartc.org/howto/lartc.ipsec.html>.

Implementing IPsec using *FreeS/WAN* requires both a kernel patch and user-space tools. Many



If our service does not support SSL, we can use SSL wrap to tunnel our plain text session over SSL.

distributions include all of the user-space tools, but the kernel may need to be patched by hand. The standard *FreeS/WAN* distribution contains the patch needed by the kernel, making it fairly easy to get it all up and running.

The configuration for *FreeS/WAN* is simple, as we simply need to define a connection and put in the appropriate IP and keying information.

An example *ipsec.conf* looks something like this:

```
# /etc/ipsec.conf - FreeS/WAN IPsec configuration file
```

```
version 2.0 # conforms to second version of ipsec.conf specification
```

```
# basic configuration
config setup
# Debug-logging controls: "none"
for (almost) none, "all" for lots.
klipsdebug=none
plutodebug=all
interfaces="ipsec0=eth1"
#interfaces="%defaultroute"
```

```
conn %default
type=tunnel
pfs=yes
keylife=2h
# How persistent to be in
(re)keying negotiations (0 means
very).
keyingtries=0
# How to authenticate gateways
authby=rsasig
# Load all connection descriptions
by default
# Some will override this with
auto=start
auto=add
```

```
conn block
auto=ignore

conn private
auto=ignore

conn private-or-clear
auto=ignore

conn clear-or-private
auto=ignore

conn clear
auto=ignore

conn packetdefault
auto=ignore

conn my-tun
left=172.16.2.100
leftnexthop=172.16.2.1
lefttrsasigkey=0sAQNmL83F
Vneuy6bQHPEtFTa3cEsq0EFvka5D
+LLSBF25L2D+8B5QIAj37GWMfL
BAg7LjrOxmV71SiID14qt3LDjtrDhS
HKVREEh6YiCqcgLMGcAfPEzlyoJ9J
z4lfxNRYbmljul4NsCgUZ9/Ob5sPxm
AG8SnEwJ0Xp1fz+SgCRYQFvX1NL
mPX5SjQU6qQQjsKMUpVfkGdsNII
BuZkm9XURhclca+YbNRJP2H3jXr
WfxZGjYttFuqjJNcNas0M16xCuTzV
VmRD3SsDSiylFsTbttZXXK4HqyHet
XIRHqjZkyMvX+JJzsgHAEo7Bw46tl
WoccmctaKSKFZR41Y8Q1gNhAvN5
wDkyntcFPO8gLVwSd4deR8F
right=172.16.5.12
rightnexthop=172.16.5.1
righttrsasigkey=0sAQPEg60Yx5u
1Bm+GHebh9MPID1ORW5sm4/q1t
pGAeQxnnCjBi75iwdKbawKA5/EPLL
h2Wrfq+qibnSQ7i16EmU8TpBYsml
V+4TUR+PYJkpTjP3tb9759EldeMon
iRCEsywsvkO4r78K1HDVw9uo07xS
l7xe2Kp+cjGt5XRHiF0nyXUCI3qkt/
niarFEI9Ih2fKfHExourF2KQvt9Lzh
A5InQF6uAncSF7M2ZUzpcsgkJG6t
dZRqQPs8TtXGR9kxSuG0pLnBcILY
4Mj5TdnTHM3JKfGXPkOzmuNH+r
DJ6tPv+DFya1GC3AfcyMsR6k7bT
OMV9MNI08qjUONCOl/DGkA/pHr2
1PfmE1Xw/etSxQxyv98/
auto=start
```

Perhaps you can tell me where I am going wrong?

Cliff Deamer, via email

As you have a minimal installation, the first question to be asked is how big your hard drive is and how it is partitioned. A functional minimal installation and a

failed complete install would indicate that you're running out of disk space, or something with your filesystems is causing problems with your install process. A full filesystem will cause the installation process to bail out halfway through and drop you back to the command prompt or shut the

system down. If you can give us some more information on your filesystem configuration, we will be better able to assist you. As you already have Windows XP installed, it may be that you do not have sufficient space on a specific filesystem for the complete install that you need.

www.linuxformat.co.uk

New Linux Desktop



Kingdom 500 Series
Broadband Media PC

"No more worms or blue screens of death it's Linux"

Features 2GHz AMD Duron
266MHz Front-Side Bus
3D/2D Graphical Accelerator
Motion Compensation
MPEG1/MPRII Video
Ethernet/Broadband Ready

17inch LCD Flat Screen
DVD/CD ReWriter+i-
80GB Ultra DMA ATA 133 IDE
DDR 266MHz memory
Wireless Keyboard and Mouse
Mandrake 10.0

Apache 2.0 Server
OpenOffice.org 1.1
K3b DVD/CD Copy
GnuCash. The GIMP
Thunderbird Anti-spam

Package 517 £549+vat (£645)

Package 519 £620+vat (£728)
with 19inch LCD Flat screen

Package 501 £310+vat (£365)
Desktop only

Extras



Wireless USB £27 +vat
ISDN PCI card £22 +vat
ADSL PCI Card £22 +vat



www.lioncom.co.uk
sales@lioncom.co.uk

tel: +44 (0)1689 889003



ANSWERS



Red hot Red Hat

Q I have a Red Hat Enterprise Server (version 2.1), with which I am having some problems. It runs on a fairly standard server. It has an Asus A7V333 motherboard and two Seagate 80GB IDE hard drive. I am using a little homegrown backup script to copy my important data from the one disk drive to the other one. Whenever I do this, it runs for a while and then the whole system locks up completely. This is not just a random thing and I can replicate it at any time. I believe it is something with the hard drive or the controller, but I'm not sure how to verify this.

Alfred, via email

A The first thing to be aware of is that as you are a Red Hat Enterprise customer you should be able to get at least installation support from them. I would also verify that all of your hardware is on their supported hardware list. Before troubleshooting any further, I would ensure that you have the newest version of the BIOS for this motherboard.

However, if you believe that this is a hard disk issue, I would suggest that you try to turn these drives down to a slower speed and see if the problem persists. First, I would recommend using `hdparm` to set everything down to safe default. Alternatively, try passing an `idebus=66` option to your kernel to force it to access the drives at a lower speed. If the problem still persists, then there is a very good chance that some of your hardware has a problem.

Server extension

Q I'm writing in the hope that you will be able to assist me in setting up *Frontpage Server Extension* on my webserver. I run a Red Hat 8 server, and although I've looked quite a lot on the web, I'm not able to get anyone's instructions to work as

expected. I'm very new to Linux and I would be extremely grateful for any guidance you could provide. Thanks for a great magazine!

Jonathan, via email

A Thanks for your email, Jonathan – let's see if we can help. In order to get *Frontpage* extensions onto your server you first need to download two packages: the *mod_frontpage* apache module and the *frontpage* server. To get a fairly recent copy of these type in the following two lines:

```
wget http://kyrian.ore.org/RPMS/apache/frontpage-5.0-1.i386.rpm
wget http://kyrian.ore.org/RPMS/apache/mod_frontpage-1.6.1.src.rpm
```

You will need to rebuild and install the *mod_frontpage* binaries from its source RPM with the following command:

```
rpm --rebuild mod_frontpage-1.6.1.src.rpm ; rpm -ivh --nodeps --force /usr/src/redhat/RPMS/i386/mod_frontpage.rpm
```

I would rename the *mod_frontpage* binaries to suit your Red Hat system by typing the following:

```
mv /usr/lib/apache/mod_frontpage.so.1.6.1 /usr/lib/apache/mod_frontpage.so
mv /usr/sbin/fpexec.1.6.1 /usr/sbin/fpexec
mv /usr/sbin/fpstatic.1.6.1 /usr/sbin/fpstatic
```

This code should end up giving you a suitable *mod_frontpage* install; however, you still require the *frontpage* package for Linux. Fortunately, this is a regular rpm installation:

```
rpm -ivh frontpage-5.0-1.i386.rpm
```

For all these modules and binaries to be of any use, you will need to make *Apache* aware of them. Open up the *Apache* config file in your favourite editor, for example:

```
vi /etc/httpd/conf/httpd.conf
```

Then, at the bottom of the

LoadModule entries add:

```
LoadModule frontpage_module
/usr/lib/apache/mod_frontpage.so
```

At the bottom of the **AddModule** entries, add:

```
AddModule mod_frontpage.c
```

The *frontpage* binary rpm includes a shell script to fix any permissions which would cause problems. To execute this script run:

```
/usr/local/frontpage/version5.0/set_default_perms.sh
```

With this complete, you should be ready to restart *Apache* by issuing a

```
service httpd restart
```

command. Now that everything is installed and running you can create webs using the following command:

```
/usr/local/frontpage/version5.0/bin/owsadm.exe -o install -p 80 -s /etc/httpd/conf/httpd.conf -xu [system username] -xg [system group] -u [FP username] -pw [FP password] -m [domain name]
```

Perl before swine

Q Hello, I've been using Linux for some years now and I've started dabbling in Perl. I'm delighted at how easy it is to use (and how forgiving it is too). As I don't work as a programmer and I don't have any heavy-duty sys-admin responsibilities I have no genuine use for it other than bettering myself and as a hobby. I've tried to get some interesting information from my system and show it on a webpage, but I keep getting an error. I thought my code was the problem first, so I've created a very original "Hello World!" script. Here is the script:

```
#!/usr/bin/perl
print "hello world\n";
```

Even this script causes the same error. My *Apache* error log shows the following:

```
Thu Feb 26 17:18:26 2004] [error] [client 192.168.1.20] malformed header from script. Bad header=Hello, world!: hello.pl
```

I'm sure it's something simple – please could you show me the light!

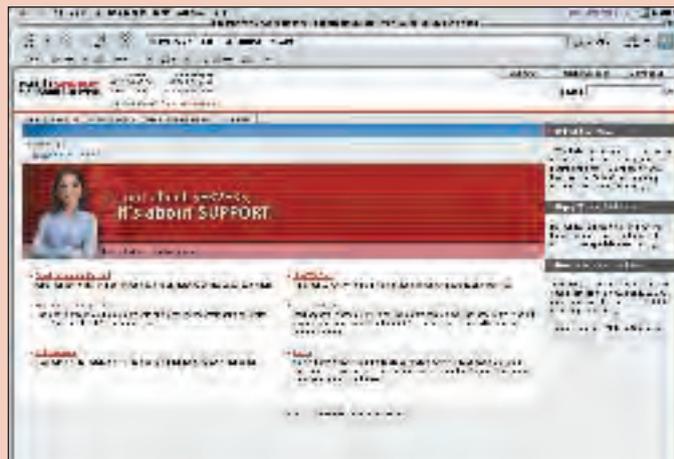
Basil, via email

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.



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!



★ Star Question – AV140 winner!

This issue's lucky winner is **Gary** – your new portable multimedia player/recorder will be with you shortly!

I am running a small server for a charitable organisation. We are using Linux to save on costs. To provide some level of redundancy the server is using software RAID level 5 over 3 disks. My question is this, what is the best way to monitor the raid set for failure? I did recently have a disk fail and as expected the system kept on running. It was only when I did a cat of /proc/mdstat that I saw the disk had failed! Is there a way to automate this so that I am notified (perhaps by email) if there is a

problem with the disks?

Gary, via email

Hi Gary, It should be relatively straightforward to write a little script to keep an eye on /proc/mdstat – as long as you know what you are looking for. Take a look at the following output:

```
# cat /proc/mdstat
Personalities : [linear] [raid0] [raid1]
[raid5]
read_ahead 1024 sectors
md1 : active raid1 sdc1[1] sdb1[1]
36840764 blocks [2/2] [UU]
```

The section, which is the most useful to us, is the last part of the

last line. The 2 U's in [U] indicate the status of the two disks in the array. If a disk is down it would replace the U for that disk with an underscore _.

A quick script, which would work for a single RAID 1 example as I've shown above, would look like this:

```
#!/bin/sh
checkraid()
{
RAIDUP=`grep -c \[UU\]
/proc/mdstat`;
if [ "$RAIDUP" = "0" ];
then
cat /proc/mdstat | mail -s
```

"RAID DRIVE DOWN"

admin@somewhere.com

```
fi
}
STAYRUNNING="YES"
while [ $STAYRUNNING == "YES" ]
do
checkraid
done
To make sure it runs properly in
the background this script should be
run with:
nohup ./script.sh &
```

A We're glad to hear your endeavours with Perl are going fairly well. Fortunately, the solution to your problem here is a fairly simple one. Because you're passing data back to a web browser, you need to tell the browser what type of data it should expect. Adding a single line as follows will tell it you're about to send HTML data and it will know to display it as that (as opposed to prompting you to save a file, etc):

```
#!/usr/bin/perl
print "Content-type: text/html\n\n";
print "hello world\n";
```

Of course, you could now insert regular HTML tags into your print statements for formatting and so on.

GRUBbing about

Q I was reading through some of my back issues (no, I have no life) and one of your previous letters regarding **LILO** and **Grub** got me thinking. I administer about 30 Red Hat servers for several clients, and with the Red Hat's change from **LILO** to **Grub**, I have no way to be certain which bootloader is being used

(without actually rebooting). Surely there must be a way to check the boot sector for this information?

A Reader, via email

A In short, yes, there is a quick and easy way to check which bootloader you are running. Assuming your hard drive is /dev/hda simply head out the first 512 bytes:

```
# head -c512 /dev/hda | strings
You should see something like this:
|lbaLILO
Df1
GRUB
Geom
```

Hard Disk
Read
Error

If you see a line containing **GRUB** a couple of lines down, then you're using *Grub*. Don't be fooled by the **LILO** in the first line, this is just an artifact of a previous *LILO* installation. For reference, a *LILO* system should look something like this:

```
|lbaLILO
DZ3q
O<.r
VVWQ
D[YSRWQ
```

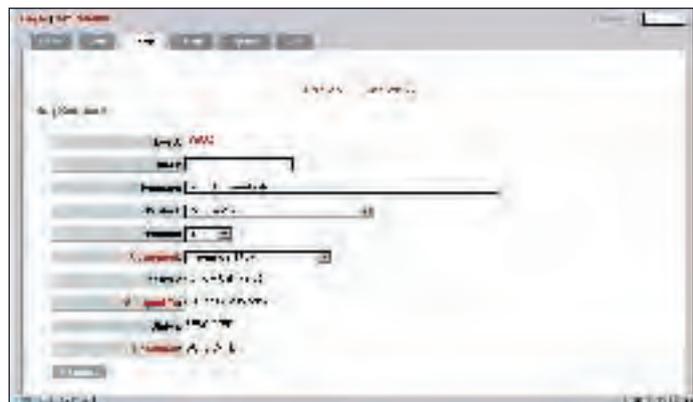
« LXFattery...

Q *LXF* is hands-down the best Linux magazine I've read, and the **ONLY** one I purchase on a consistent basis!

I installed Fedora from the disc that came with *LXF49*, and so far it has been great but for a few things. My initial install was in a dual-boot environment, in which I installed over Red Hat 9 (an overwrite, not an upgrade). Afterwards, I was easily able to download my updates and retrieve my personal data from storage media. In fact, I had Fedora running as smoothly as RH9 had been within about an hour. Then, after numerous frustrations with Windows XP, I decided to scratch the dual-boot idea and install FC1 over the entire hard drive. The installation went fine, but after post-setup, I found that I couldn't use *up2date* without the program hanging at some point, either at the 'Fetching RPM Header' stage, the download stage or the installation stage. Suffice to say, it's frozen at about every point, and I've run it enough time to make my head spin. Apparently, this is common, as I've found some postings pertaining to it, but none of the info provided in them has worked. Can you suggest some (hopefully simple) steps to alleviate this problem?

Also, just another small grievance with the OS: sometimes I can't seem to unmount my CDs or DVDs. I have a Samsung CD-ROM SC-140C and a Sony DVD-RW DW-U12A. This problem applies to both drives and with all types of media; DVD movies, DVD-/+R(W) and CD-/+R(W). Each time this happens, I always make sure to close any *Nautilus* windows pointing to the drives, shut down any apps using them, and often I will wait up to ten minutes after this and will still have no luck unmounting them. Even unmounting as root from a terminal has no effect, and this is becoming a daily problem. Hopefully you've heard of something like this before, because I'm totally baffled. To end this letter on a less-whiny note, I'd like to thank you especially for the monthly DVD.

It's a nice treat to get free software, and even better when it's GOOD software and LOTS of it at that! You guys are alright!
Robert Pierce, via email



Red Hat's Bugzilla is a great resource for users trying to solve problems with Fedora Core or Red Hat's earlier distributions.

A Thanks for your kind words! Looking at the Bugzilla for Red Hat, there is a known issue with *up2date* if it is unable to fetch the required files from the FTP server. Rather than gracefully handling the problem, the front end simply gets 'stuck' and doesn't do anything useful. You can read about the problem at http://bugzilla.redhat.com/bugzilla/show_bug.cgi?id=88349. It sounds like it is an intermittent error, so you might try it five times and never see the problem, then do it once again and it will occur. Sadly, until they fix the specific servers which are causing the error, it will continue to be a problem unless the code is fixed.

Your CD/DVD unmounting problem is yet another known bug in Fedora Core, which is referenced at http://bugzilla.redhat.com/bugzilla/show_bug.cgi?id=109687. It would appear that if you kill the *fam* process, you will be able to **umount** your drive correctly. As Fedora Core is still in beta testing, little quirks like these are to be expected, as annoying as they are.

Slacker

Q I've got several computers that I want to install Slackware onto. One in particular is a Dell Latitude CPX laptop that only likes Slackware. I've got the *LXF48* DVD, but rather than burn more CDs or cart the one DVD around all the computers that I want to install it onto, setting up a NFS server and getting the other computers to install from that would be ideal.

In the 'select source media' dialog, there is an option for installing from NFS, but I get stuck when it asks for the location of disk *a3*, which appears to date back to

when Slackware came on a floppy set. Trying to manually mount NFS directories locks up the console that I'm on, and I've run out of ideas – please help!

Richard Smith, via email

A When doing a NFS install, the installer needs to know where to find the various disk images for the installation. If on the remote server, the *a3* image is located in */mnt/slackware/a3*, then you need to pass */mnt/slackware/a3* to the installer. Why the manually mounting of the filesystem causes the console to lock up is worrying, but likely down to a network or DNS issue.

If you are interested in trying another distribution onto your Dell, check out <http://linux-laptop.net/> and search for the model you have. People have documented their installs of Linux onto many different systems, so if you are trying to make Red Hat, Fedora or Debian work with it, then you may find some tips and hints to help you along the way.

Mozilla moan

Q I am having problems with Mozilla (I think)... When I try to connect to the Internet/find a website, the following message is displayed: HTTP Status 500 - No Context configured to process this request type Status report message No Context configured to process this request description. The server encountered an internal error (No Context configured to process this request) that prevented it from fulfilling this request.

This leaves me none the wiser. I contacted my ISP's 'technical' helpline, and they confirmed that their servers are operating

correctly. When I told them what software I was using, they effectively told me to bugger off. I am pretty sure the problem lies within *Mozilla*, because when I open History, and click on a previously visited site, I am immediately connected! I seem to be connecting to the server OK – *wvdial* goes through the negotiations etc), but the above error message appears in the *Mozilla* window.

Tony Bradley, via email

A The first step to debug this problem is to disable any proxy settings within *Mozilla*, and attempt to open a website that you are completely confident about being up – <http://slashdot.org/> for example. It's important to limit the number of servers involved, otherwise we won't know if it's the proxy or remote server. A 500 error is generated by the server, although it's likely that it's a malformed request which causes the server to generate such an error.

If you believe this to be a *Mozilla* problem, copying *~/mozilla* out of the way and reinstalling with the current build of *Mozilla* should be a reliable way to test the system. **LXF**

Submission advice

We are happy to answer all sorts of Linux-related questions. If we don't know the answer, we'll find out for you! But in order to give you the best service, it helps a lot if you read the following submission advice.

- Please be sure to include any relevant details of your system. "I can't get X to work" doesn't really mean anything to us if we don't know things like what version of X you are trying to run, what hardware you are running on.
- Be specific about your problem. Things like "it doesn't work" or "I get an error" aren't all that helpful. In what way does something not work? What were you expecting to happen? What does the error message actually say?
- Please remember that the people who write this magazine are NOT the authors or developers of Linux, any particular package or distro. Sometimes the people responsible for software have more information available on websites etc. Try reading the documentation!

We will try 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 really give personal replies to all your questions.

WRITE TO US AT:
Linux Format, Future Publishing, 30 Monmouth Street, Bath BA1 2BW or email: linuxformat@futurenet.co.uk

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).



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



February 2004

Product code:
LXFB0050(cd)
LXFD0050(dvd)

DVD HIGHLIGHTS:

EVERY SINGLE HOTPICK EVER! All the programs that have featured in LXF's Hot Picks selection from issues 1 to 50

CDs HIGHLIGHTS:

50 HOTTEST PICKS! Including: Anjuta, ZNES, Planner, Scribus, Mozilla, Anjuta, Exim, ProFTPD...

MAGAZINE FEATURING:

Big Hitters biog collection, Extreme Programming Pros & Cons, FORTRAN, CrossOver Office, Systemax 6507 & Evesham Reliance servers, Astronomy apps, Blender



January 2004

Product code:
LXFB0049(cd)
LXFD0049(dvd)

DVD HIGHLIGHTS:

Pathatizer, Kopete, LDP, K3b, AbiWord, Scribus, CyberCafeOrganizer, Warewolf, Ant eater

CDs HIGHLIGHTS:

Fedora, Vega Strike, Omnis Studio 3.1, Mandrake 9.2 updates, OpenOffice.org Quickstarter, ReCall, The GIMP, NetPanzer

MAGAZINE FEATURING:

Java Desktop Environment, Linux DVD Players, CD copy protection feature, Maya 5, Smoothwall Corp, Guardian, What On Earth is the AFFS? RH Enterprise Linux 3



Christmas 2003

Product code:
LXFB0048(cd)
LXFD0048(dvd)

DVD HIGHLIGHTS:

KDE3.2 alpha2, DamnSmallLinux, Audacity, Samba 3.0

CDs HIGHLIGHTS:

SUSE Live Eval 9.0 (boots from CD/ISO), Slackware 9.1, gentoo File Manager, Karchiver, BitTorrent, Mozilla, Opera, Firebird, Mamba Open Source Site Server

MAGAZINE FEATURING:

OpenOffice.org & StarOffice, HP Itanium server, KDE3.2, MIDI sequencers, BitTorrent, Majesty review, Scribus DTP tutorial, Slackware 9.1 review, Gigabyte SR125ED review



December 2003

Product code:
LXFB0047(cd)
LXFD0047(dvd)

DVD HIGHLIGHTS:

OpenOffice.org (PPC), MailScanner, XFCE, LDP, PuppyLinux, Slurm, flphoto, Battlemech, XMMS, Fandango, CVSToys, Scrubber, Sylpheed

CDs HIGHLIGHTS:

Mandrake 9.2, OpenOffice.org, GNOME 2.4

MAGAZINE FEATURING:

Mandrake 9.2, Linux Expo UK report, Lazarus, CVS, Xinit SharpStation, MTStudio, SUSE 9.0 review, Scanners tutorial, Newsreaders roundup, IBM's Linux strategy, Nessus

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
FORMAT

www.linuxformat.co.uk

PRIORITY ORDER FORM:

I would like to subscribe/renew to *Linux Format* (please delete as appropriate):
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

 Media with passion	Instruction to your Bank or Building Society to pay Direct Debits	
	 DIRECT Debit	Originator's Identification Number 7 6 8 1 9 5
Future Publishing Ltd, Cary Court, Somerton, Somerset, TA11 6BR	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)

4. Bank or Building Society account number _____

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) _____

Banks and Building Societies may not accept Direct Debit Instructions for some types of account.

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 12.05.04

OFFER CODE LFXFP53

LINUX

FORMAT

» 3 EASY WAYS TO ORDER:

1 PHONE 0870 444 8645

2 POST Future Publishing Ltd, Freepost, TA11 6BR

3 FAX this completed form to: 01458 271146



OVERSEAS CUSTOMERS CALL:
+44 870 444 8645

SUBSCRIBE AND SAVE!

SAVE 29% OFF THE COVER PRICE

That's 13 issues for the price of 10
Don't miss out, call today

- ✓ Save over **£24*** on the price you'd pay in the shops
- ✓ **FREE** delivery direct to your door
- ✓ **Never miss an issue!**

LINUX FORMAT BRINGS YOU:

- 1** Essential reading for all Linux users
- 2** 20-page *Linux Pro* supplement with every issue
- 3** Three CDs or a DVD packed with Linux software
- 4** Latest news, latest reviews and the best Linux tutorials around

Taking out a subscription or giving a gift is easy. Simply call **0870 444 8645** or complete the order form opposite.

*£24 saving based on price of DVD issue.

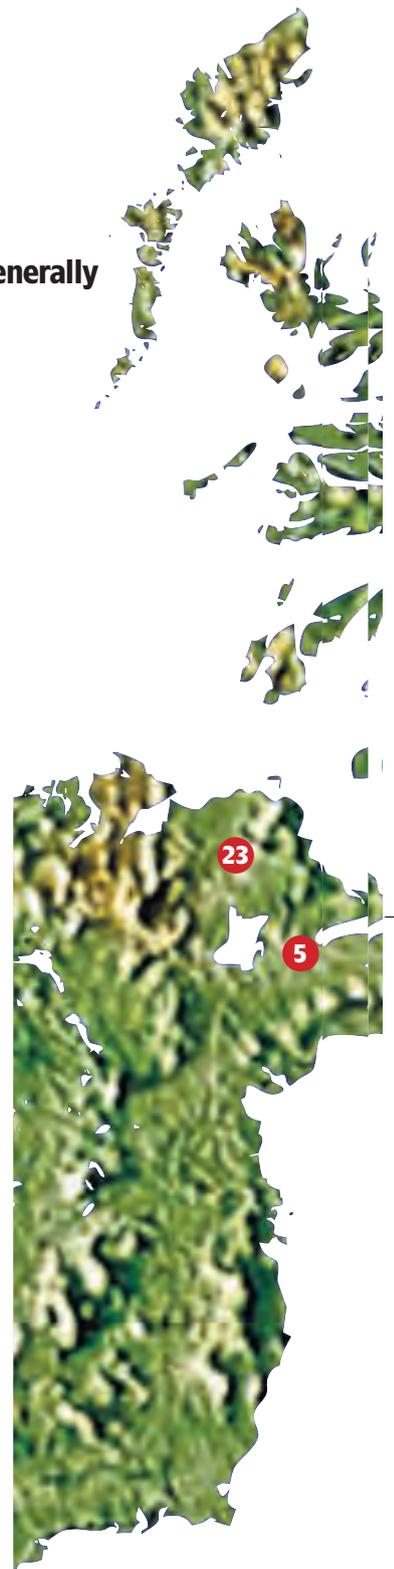


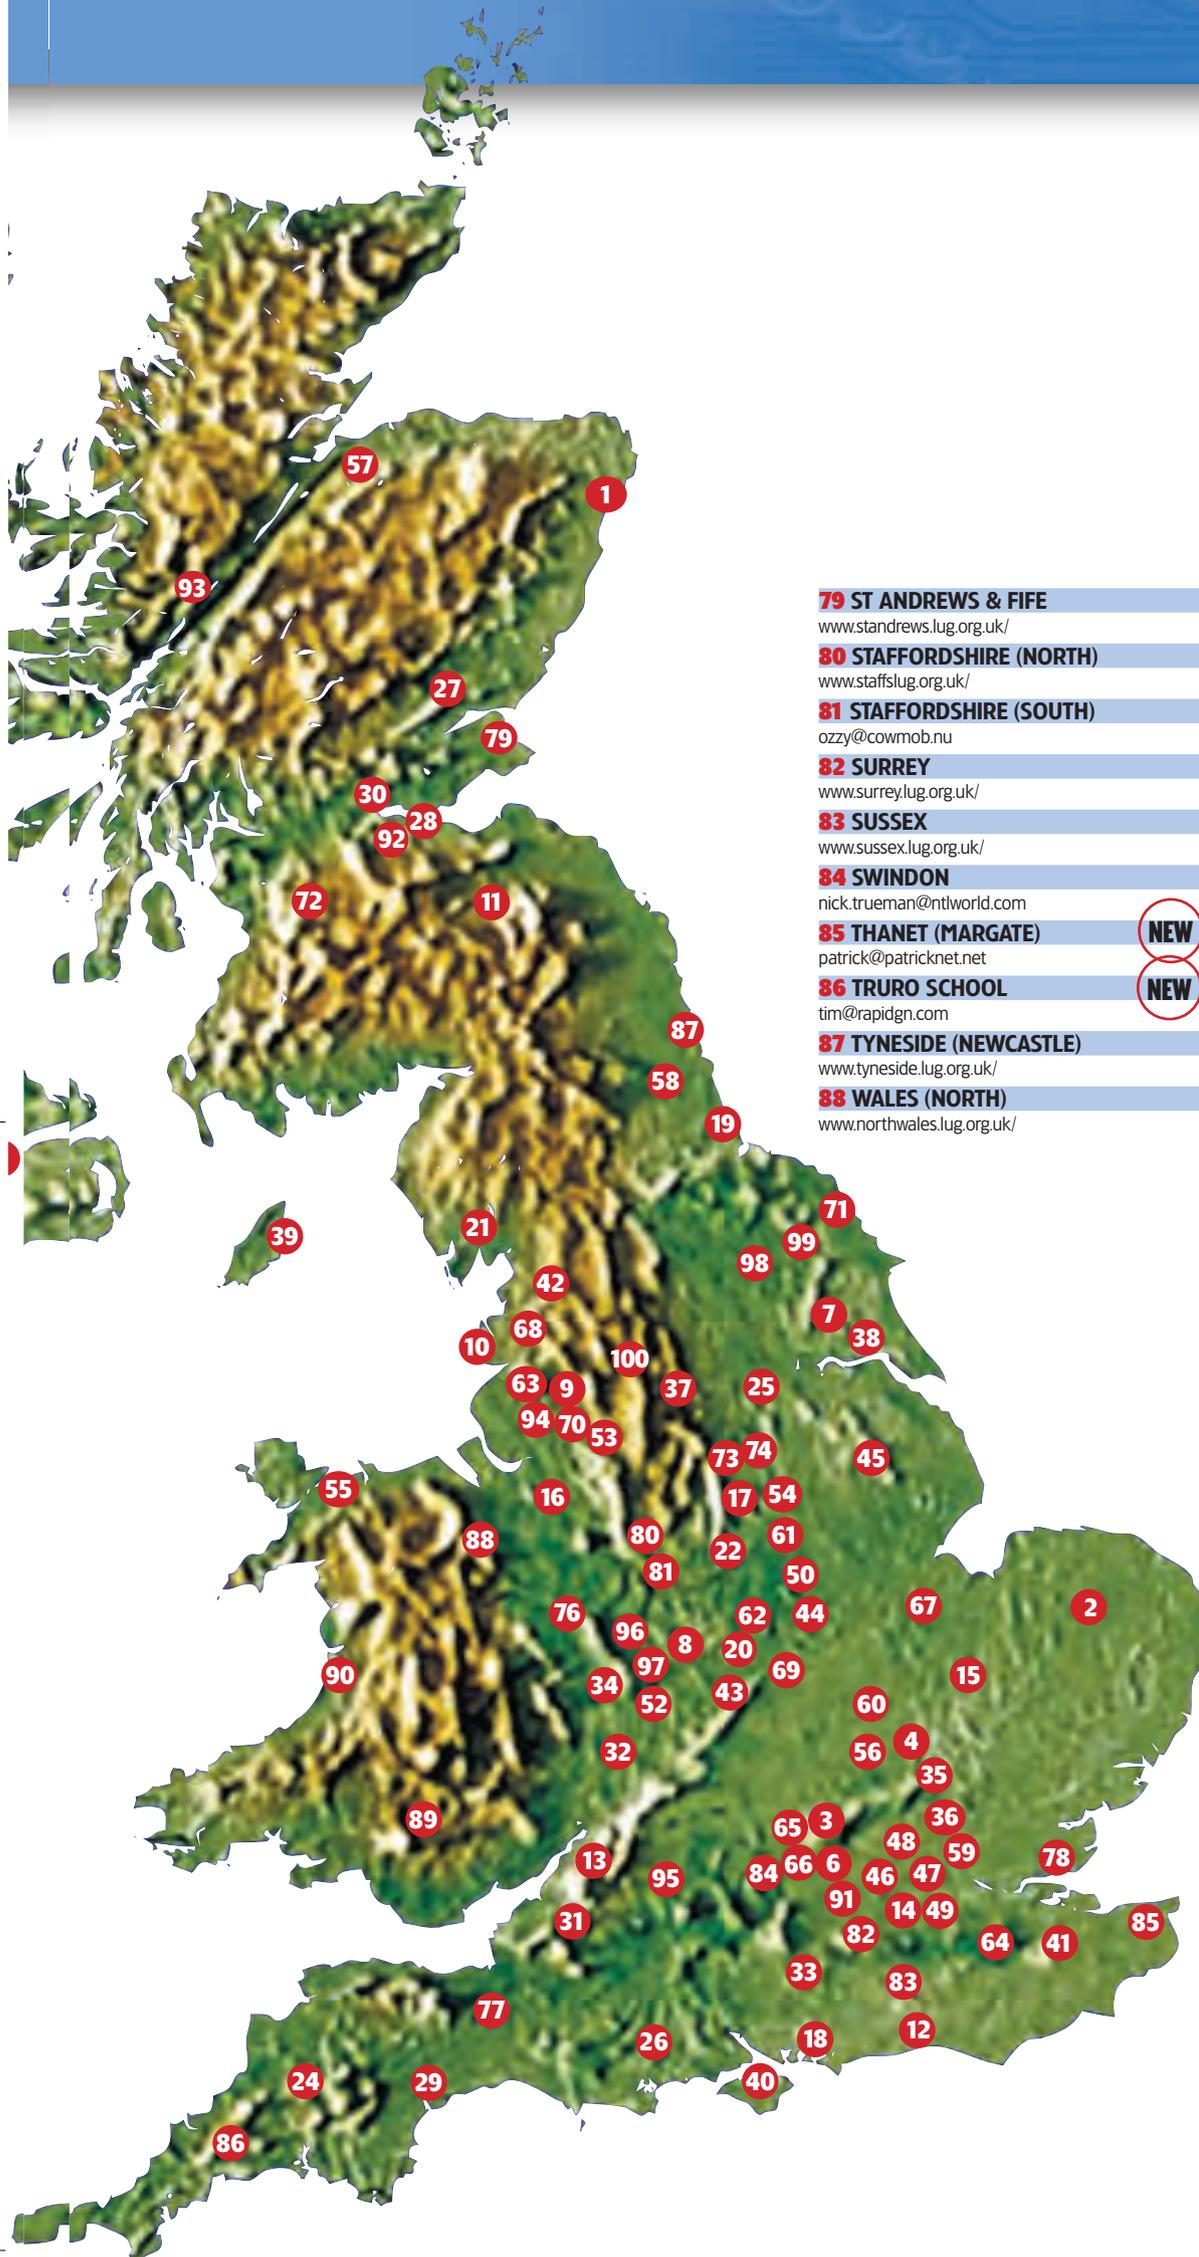
ORDER HOTLINE ☎ 0870 444 8645

User Groups

LUGs worldwide are full of members keen to help with your problems, discuss ideas, and generally natter about all things Linux. You can find lots more information online at www.lug.org.uk

1 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 NEW www.blackpool.lug.org.uk/	36 HERTFORDSHIRE (EAST) www.earlug.org.uk/	62 NUNEATON www.nuneaton.lug.org.uk/
11 BORDERS (ENGLAND/SCOTLAND) welby@mcrobbers.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 NEW ian@siim.info	69 RUGBY NEW www.rugby.lug.org.uk/
18 CHICHESTER HARBOUR - MARITIME GEEKS NEW 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) NEW 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



**79 ST ANDREWS & FIFE**www.standrews.lug.org.uk/**80 STAFFORDSHIRE (NORTH)**www.staffslug.org.uk/**81 STAFFORDSHIRE (SOUTH)**ozzy@cowmob.nu**82 SURREY**www.surrey.lug.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/

LINUX USER GROUPS

LUG OF THE MONTH

UK UNIX User Group (UKUUG)

UKUUG was formed to represent users of Unix and Open Systems in the UK. Complete funding by membership subscriptions allows total independence from specific hardware and software vendors, and any profits are used to further the activities of the organisation. UKUUG's importance is underlined by the fact that every major supplier of Unix, Unix-related and Open Systems is now a member of the group. UKUUG holds regular events for members, and occasional events for the general public.

We have two main conferences a year; the Winter Conference in or around February, and the UK's main Linux Conference in the summer.

The **2004 UK Linux Developers' Conference** will take place in **Leeds** from **Friday 30 July to Sunday 1 August**, preceded by tutorials on Thursday 29 July. Topics presented in recent years have included: ARM Linux, Benchmarking, Clustering, CORBA, Debian Package Management, Enterprise Filesystems, *Exim*, *Flightgear*,

GNOME, *Heartbeat*, I2O, JFFS, KDE, Mail Servers, Memory Management, Performance Programming, *Powertweak*, *Pymmetry*, *Samba*, Security, SMP, *VMware*, *Zerocopy*, *Zope*. More information about this and previous events is available on our website at www.ukuug.org/.

A broad cross-section of the Linux community will be gathering for the conference: it's a great way to widen your knowledge and keep up-to-date with what's happening in the world of

Linux. This low-cost event is for anyone with a serious interest in Linux, including systems administrators, Linux pros, developers and enthusiasts from companies or Linux User Groups throughout the UK and beyond.

To keep the conference fees low, we are seeking sponsors and exhibitors. For further information about sponsoring, exhibiting, or attending the event please contact the UKUUG office by emailing office@ukuug.org or by telephoning +44 (0) 1763 273475

WORLDWIDE LINUX USER GROUPS

Ideas needed from global Free/Open Source Software users

Our brand-new UK map this month marks the beginning of *Linux Format's* efforts to get more involved generally with Linux User Groups communities, be it in the UK or worldwide. We have one small problem though: we can't actually agree on how we can do this! Following last issue's very brief appeal at the end of *The Great 2004 LUG Explosion*, we've had several good ideas already:

■ Mumbai reader **Vishal Kumar** suggested that some worldwide LUGs – particularly in Asian countries like Thailand where Linux adoption is being fostered by local governments – might be able to benefit from having the opportunity to be subscriptions agents and/or distributors of our magazine in any areas of the world where *Linux Format* isn't readily available on the news-stand. We're exploring ways of making this happen as you read this: if your LUG would be interested in taking part in such a scheme, please do get in touch as soon as possible – we'd like to hear from you so we can get a very rough idea of numbers.

■ **Adler Milstrey** emailed us from Berlin with the suggestions that user groups in wealthier states could practice some sort of 'twinning' scheme with LUGs in developing areas. If you've got any ideas as to how *Linux Format* could help facilitate this, let's hear them!

■ **Dougal Greenwood** from New South Wales, Australia maintains that a major factor in holding back Linux adoption for many computer users is that they aren't aware of the fact that you can run Windows and Linux side-by-side. He suggested that a good way of publicising the existence of LUGs to people in your local area could be to co-ordinate a 'Windows Amnesty Day', whereby members of the public bring their PCs along and are helped through making their own setup dual-boot with both Windows and Linux by LUG gurus.

LXF reckons that this idea gives some much-needed branding and direction to the otherwise limited appeal of the 'Linux Install Day' idea – if maybe this were organised on a national or international footing, and perhaps sponsored by Linux software vendors, we suspect that publicity would be quite easy...

So now it's down to you! Got an event happening? Need another LUG to send you some distro CDs from overseas? If you don't send details in, we can't publicise them! If you have any ideas about what *Linux Format* can do to help LUGs, eg suggestions for topics, regular sections or events to be covered on these pages, you don't have to be a LUG member to contribute! Email us at lxf.lugs@futurenet.co.uk or if enough of you request a new LUG Forum on our website, we'll start one if there is sufficient interest!



Jono Bacon believes that personally putting Linux in the face of the computer-using community is the way forward...

As I mention in my Comment column this month, advocacy in its truest form is all about getting people already sat in front of a computer to give Linux a try. The concept of 'seeing is believing' should be the rule of thumb here. If you say the words 'Linux', 'Free Software' and 'Open', many people will conjure up thoughts of complexity, unfamiliarity, geeks, and maybe even piracy if they don't understand *how* Linux is 'free'.

These preconceptions can only be levelled out by getting someone in front of a Linux machine. The problem here of course is actually getting hold of a PC in a place where people are located. There are however a few options. One simple method is to always carry a Knoppix CD with you when you are out and about. If you are in a computer shop and get talking to someone about Linux, or if you are at a friend's house, you can boot their computer into Linux and give them a tour.

With the price of CDRs being so low, it is also viable to send a stack of bootable CDs to places such as schools, community groups, charities and other

areas. Another useful option is to post some CDs through your neighbours' letter boxes – if this business practice works for AOL, why can't it work for Linux? This can be particularly worthwhile if your neighbours are likely to be regular computer users who want to write documents, browse the web, send email etc. One caveat though: if you are successful at this sort of rather evangelical advocacy, you may end up getting a lot of knocks on your door asking for help. If this is the case, you can kindly suggest the typical resources available if you don't have the time to help everyone personally.

This Advocacy column is taking a break for a few months; in the meantime, remember that *practical* advocacy is key to getting the message out to people, charities and businesses.

Don't forget to check out the new Business LUG site featured in *LXF51's* User Groups (www.businesslug.org); and as ever, please send your comments, advocacy stories and suggestions to spreadingtheword@jonobacon.org 

Linux User Group organisers

If you're not listed here, or we have your details wrong, please contact us at: **LUGS!, Linux Format, 30 Monmouth Street, Bath, BA1 2BW** or email your details to: lxf.lugs@futurenet.co.uk

Coverdisc

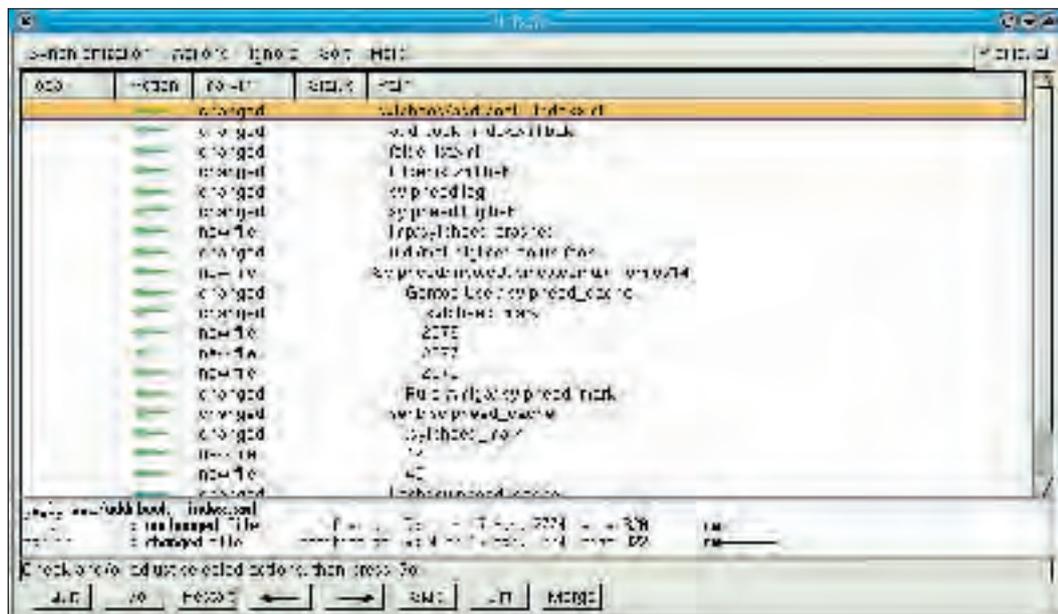


Neil Bothwick is your guide through the wonders of this month's jam-packed *Linux Format* DVD. There's MandrakeClub info here pertinent to CD readers too!

DESKTOP UNISON

Many people now use more than one computer. Whether it be one desktop at home and another at work, a desktop and a laptop or a combination of these, sooner or later you'll find yourself using one computer and needing files from another. There are various ways of dealing with this, you could copy files to a floppy disk or USB flash drive, or you could use rsync if the two computers are networked. Each of these solutions has limitations, especially if files have changed on both computers since they were last synchronised.

Unison is intended to provide easy two-way synchronisation of



Nothing at all to do with trade unions, *Unison* keeps your files in sync on different computers.

directories on two computers (or even on one). It will compare two directories (or two lists of directories if you set up a suitable profile file) and show you the files to be copied in either direction. It keeps track of the previously synchronised state, so it can cope with the situation where the same file has been updated on both computers since they were last synchronised, rather than simply assuming that the newer version of any file is all you want to keep.

It will synchronise two directories on the same computer, so you could sync your home directory on your work computer to a USB flash disk, then sync this with your home computer. Unison works over any TCP connection, so it is also possible to synchronise computers over the Internet. It uses the rsync algorithm when copying files. This only transfers the changes to a file rather than the whole file, making it suitable for use on low bandwidth connections, and faster on any sort of connection.

Unison runs on Windows as well as Linux, so you can still transfer files

from work, even if your employer has a less enlightened choice of operating system. There are graphical and text interfaces, the former being the easiest to use, the latter being ideal for using in scripts. The DVD contains statically compiled binaries if you want to quickly try it out, as well as the usual source code, and comprehensive documentation.

GAMES UT2004DEMO

You can read a review of *UT2004* elsewhere in this month's *Linux Format*, and now you can try it for yourself with the demo version included on the DVD. Because of the amount of space needed for Mandrake 10.0 and the other programs on the DVD, we had to include the UT2004demo installer as a compressed file, so installation requires a few steps. First, copy the file `ut2004-lnx-demo-3120.run.bz2` from the `Games/UT2004demo` directory of the DVD to your home directory and unpack it with

```
bunzip2 ut2004-lnx-demo-3120.run.bz2
```

Now run the installer with

```
sh ut2004-lnx-demo-3120.run
```

By installing it as a normal user, you keep all the files in your home directory, making the installation cleaner, and easier to remove later should you wish to (maybe because you have bought the full game). If you need to make the game available to more than one user, you can install as root, in which case it will be installed into `/usr/local/games`.

Installing as a user creates one minor problem with the default settings that is common to all games that use the Loki installer. It tries to create a link to the program that has the same name as the default directory name. To avoid this, make a slight change to the name of the installation directory, say from `~/ut2004demo` to `~/ut2004-demo`.

Now you can run the game with

```
./ut2004demo
```

There is an uninstall script within the installed directory, so you can remove the demo with

```
sh ut2004-demo/uninstall
```

On the DVD



Wherever you see this logo it means there's related stuff on the DVD

IMPORTANT NOTICE

Before you even put the DVD in your drive, please make sure you read, understand and agree to the following: The *Linux Format* DVD is thoroughly tested for all known viruses, and is independently certified virus-free before duplication. We recommend that you always run a reliable and up-to-date virus-checker on ANY new software. While every care is taken in the selection, testing and installation of DVD software, Future Publishing can accept no responsibility for disruption and/or loss to your data or your computer system which may occur while using this disc, the programs or the data on it. You are strongly advised to have up-to-date, verified backups of all important files. Please read individual licences for usage terms.

CAN'T BOOT, WON'T BOOT?

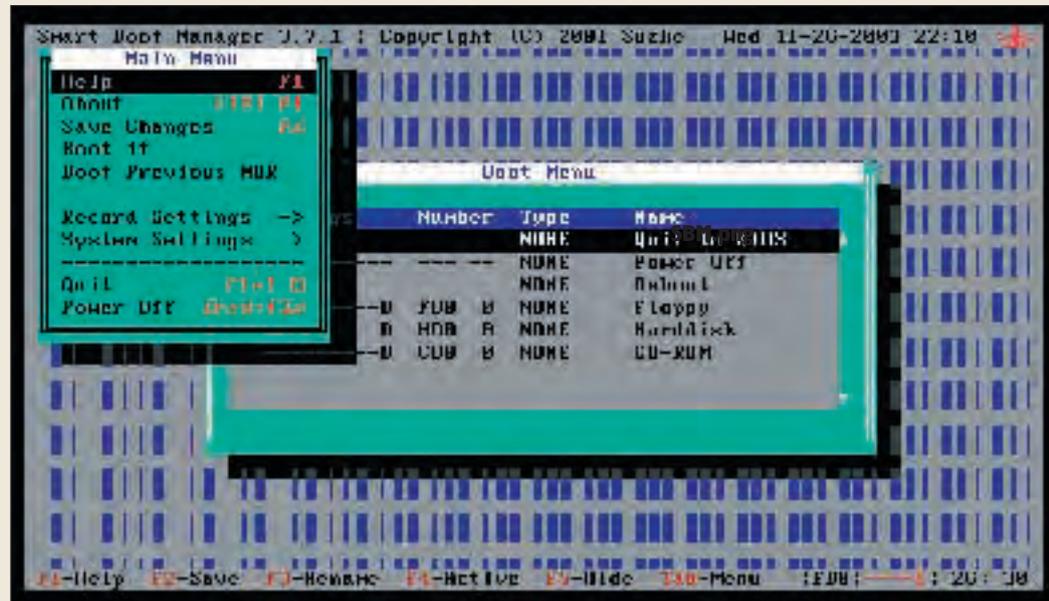
Troubleshooting your LXF coverdisc operation

Although many of our coverdiscs – including this month's – are made to be bootable, from the emails we've received, it seems that a few people have experienced problems in trying to boot from them.

If this problem 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 manual should give details on 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 particular 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 *isohybrid* 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



Use *Smart Boot Manager* to persuade a reluctant computer to boot from your *Linux Format* DVD.

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

sbootmgr.iso. Burn this to a CD-R and boot from that. When you see the boot menu, eject the CD and replace it with the coverdisc before selecting the relevant device (usually CDO).

If the disc boots, but then the system hangs before the installer comes up, try

unplugging any non-essential USB devices from the computer. USB keyboards and mice are fine, but some scanners and cameras have caused problems in the past. You should still be able to use them once the system is installed – just leave them off for now.



You have read the review, now play the game. This is a playable demo of *UT2004*.



COVERDISC DVD

CREATING INSTALLATION CDS FROM THE DVD

Under Windows or Linux – the choice is yours!

As is now normal practice with *Linux Format's* 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. For convenience, this can be done from either 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, type the following command in a terminal `sh /mnt/cdrom/Distros/Mandrake/mkiso` This will create the three 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/Mandrake/mkiso /tmp/iso`

Note that you should not `cd` to the directory on the DVD when running this script. It will not fail as it used to on previous *LXF* discs, but it will be slower because the script is not able to open the cache file used to speed up the creation of the second and third discs.

If you are short of space, it is now possible to create single ISO images with `sh /mnt/cdrom/Distros/Mandrake/mkiso -d 1` or `sh /mnt/cdrom/Distros/Mandrake/mkiso -d 2 /tmp/iso`

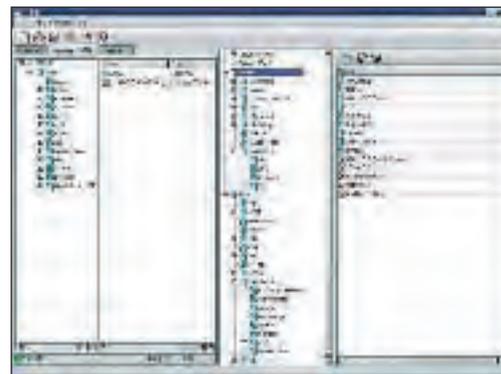
If you get an error message along the lines of `../Essentials/jigdo/jigdo-file: Permission denied`, it means your DVD has been mounted with the `noexec` option, which prevents running programs directly from the disc. For security, this is implicit if the DVD is mounted with the `user` option. To fix this temporarily, type this as root:

```
mount /mnt/cdrom -o exec,remount
```

Alternatively, you could install *jigdo* to your computer, from the Essentials directory, and use the *mkiso2* script. This should use any copy of *jigdo* on your hard disk in preference to the DVD version, which will avoid this problem.

The *mkiso2* script has not been tested on many computers, so please let us know of any problems you have with it.

After creating ISO images with the *mkiso* script, you can burn them to CD with *K3b*, which just happens to be on this month's *LXF* DVD.



Creating ISO images in Windows

Windows users can create the CDs by double-clicking the *winmkiso* icon or running the script from a MS-DOS prompt. In the latter case, you should change to the Mandrake 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, this should be effective:

```
E:
cd Distros\Mandrake
winmkiso D:\ISO
```

DVD CONTENTS AT A GLANCE

Desktop

CDRDAO	Disk-At-Once Recording of Audio CD-Rs
CDRtools	A tool to create disk-at-once and track-at-once CDs
Cpre	A regex-enhanced supplement to cp
DuracellCPULoadMonitor	CPU load monitor using an AA-battery tester as its display
Eggcups	A print-queue panel icon
EvRouter	A Linux input event router
GRegExpExplorer	A graphical regular expression explorer based on PCRE
GWhere	A CD catalog manager
IFtop	Real-time bandwidth usage information
K3b	A CD/DVD burning application
KAlarm	An alarm message, command, and email scheduler for KDE
KDar	A disk-based graphical backup and archiving program for KDE
KDENetworkMapper	A network discovery and analysis tool for KDE
KKeyled	A KDE applet that displays the keyboard lock states
Komics	A KDE Applet that fetches comics from the Web
Konserve	A small backup application
MultiMediaKbdController	A project to enable mouse gestures and multimedia buttons
PWSafe	A command-line password database
Revelation	A password manager for GNOME 2
Scorcher	Automate common CD/DVD rip, encode, and burn tasks
TKsync	A GUI to rsync
Type1URWFontsWithCyrillics	Type1 URW fonts extended with Cyrillic glyphs
TypingTrainer	Software to exercise typing skills
Unison	File synchronizer
X10UniversalDeviceDriver	A device driver for X10 home automation for script access
XFileExplorer	A file manager for the X Window System

Development

AccessModifierPlug-in	Change the visibility of Java elements in Eclipse
ADODB	A PHP database abstraction layer
Bitmask	A light-weight and fast 2D collision detection library
CDInputAndControlLibrary	A CD-ROM read and control library with an ISO 9660 library
Flyspray	An easy to use, Web-based bug tracking system
GCC	The GNU Compiler Collection
gPHPEdit	A GNOME2 editor for editing PHP and supporting files
Jagzilla	A Java-based set of components for Bugzilla
Loudmouth	A Jabber client library
PkgBuilder	A software package system for Linux, similar to BSD ports
Pylint	A lint tool for Python
Qt	A GUI software toolkit
RunTimeAccess	Gives easy run time access to your program's data
SpringUI	Eclipse plugins to ease working with the Spring Framework

Distros

Mandrake	The latest release of this popular Linux distribution
----------	---

Games

Concentration	A memory game
HuntForGold	A real-time strategy game in Java
LiquidWar	Eat the other player's fighters
SolarWolf	An action/arcade game written in Python and Pygame
UFO-AlienInvasion	A strategy game inspired by the X-COM series
UT2004demo	Multiplayer first person shooter

Graphics

ETVPersonalVideoRecorder	Personal video recorder (PVR) especially for European use
FFTV	A viewer and recorder for TV and radio
File2DIVX3pass	Mencoder frontend to convert any media-file to Divx
Fotoon	A tool to organize photos in an easy and fast way
GQview	An image browser and viewer
GrabMV	A tool that acquires microMV video from digital camcorders
KDEImageDatabase	Index/browse/search large numbers of images
KolourPaint	An easy-to-use paint program for KDE
MVC	A video capture tool with motion detection
Recordshow	A script to record radio/TV from a Linux V4L device
RivaTV	A Video4Linux kernel driver for nVidia-based graphics cards
Scale2x	An image enhancement graphics effect
VcdPyCoder	A graphical front-end to video conversion tools
VideoEncodingScriptCreator	A GUI for easily creating video encoding scripts
Vobcopy	Copies DVD .vob files to harddisk
Xvidcap	A desktop video screen capture application

Help

LDP	A complete mirror of the Linux Documentation Project
-----	--

Internet

AlvarosMessenger	Multilanguage MSN Messenger client
Balsa	GNOME mail client
DC-QT	A Qt GUI for the dctl Direct Connect client
EasyMoblog	For the publishing of blogs with updates via email
EPIC4	An ANSI-capable textmode IRC client
FreeWorld-IRCD	The FreeWorld IRC Daemon
Gossip	A GNOME instant messaging client
HarvestMan	A full-featured, multi-threaded offline Web browser
JunkMailBufferingAgent	A tool to queue email messages until the sender is verified
Kjrm	KDE port of Yahoo Messenger
LightHawk	A Jabber client for KDE
Lurker	A full-featured email archiver and search engine
MozillaThunderbird	A total redesign of the Mozilla mail component
POP3spam	A tool that deletes unsolicited mail from a POP3 server
Popselect	A tool to manage a POP3 mailbox

JOIN THE CLUB

www.mandrakeclub.com

When you like a distro and want to give something back, a typical approach is to buy a boxed set, thereby rewarding the producers and coders for their work. Sadly, a large chunk of the cost of a boxed set goes on duplicating discs, printing manuals, producing packaging, shipping and margins for distributors and retailers – only a small proportion goes to the people who created the distro. MandrakeSoft's answer to this – that works alongside its boxed sets – is MandrakeClub. By joining the club, almost all of your subscription goes to Mandrake, the only middlemen taking a cut are the credit card companies. Club Membership benefits include:

- ★ Full access to commercial applications that are normally only available in retail packs. Over 130 high-quality packages built and tested for Mandrake Linux can be installed with a single mouse-click.
- ★ Special download mirror list exclusively for MandrakeClub members – you can download extra CD images, even the full *Power Pack* distribution, some commercial packages and extra documentation

- ★ Special discounts on third-party products and services. Past promotions have included TransGaming's *Kohan* games, *Hancom Office*, Yopy PDAs, Win4Lin, and many others
- ★ Special discounts for products and services at MandrakeStore
- ★ Direct-trading program for buying and selling MandrakeSoft stock, plus access to real-time trading information
- ★ Multilingual discussion forums
- ★ Multilingual articles
- ★ Vote on which packages should go into future distributions
- ★ Access to MandrakeClub.com – a place where your voice will be heard plus access to special 'club-only' articles
- ★ A huge database of all the software packages that are currently available for the Mandrake Linux distribution. More than 50,000 RPM packages – over 100GB of software (and still growing) – is now available for all versions of Mandrake Linux!
- ★ Full access to MandrakeSoft's e-learning platform, which provides



MandrakeClub offers many benefits, including extra software and that warm fuzzy feeling of knowing you are helping. Here you can vote for and test packages that you would like to see included in the next release.

comprehensive courses about Linux and Open Source software, plus quizzes to validate your knowledge

- ★ Early access to Mandrake Linux manuals and technical documentation

If you are interested in getting extra features, while putting something back into the creation of Mandrake Linux, you should take a look at what's on offer at www.mandrakeclub.com. [LXF](#)

QuickFile
Sylpheed
TorrentFlux
UrlGet

Thunderbird extension to add keyboard-only message filing
A GTK+ based user-friendly email client
A PHP Torrent client and manager
A download manager with a GUI

Office

Group-Office
PTimeTracker
STORM

A Web-based groupware office suite
Track the time spent on different activities
Web-based tools to help IT management in schools

Server

ApacheTop
Eucalyptt
Fakepop
GiftWeb
glFtpD
HordeApplicationFramework
IMP
iWare
Jigsaw
LeahHTTPDaemon
MailScanner
Mnemo
mod_log_sql
MySQLApache
NullWebmail
PhilsWebmail
phpPgAdmin
PunBB
Sioux
WebApplicationGateway
Wzdtftp

A curses-based top-like display for Apache
Open discussion forum with simple moderation
A fake POP3 server which never has any mail
An online gift registry system
FTP Daemon for Linux. Great program for an ISP or anyone!
A PHP application framework
A Web-mail system based on IMAP and PHP
A Web-based content management system
W3C's leading-edge Web server platform
A secure, stable, fast, and light Web server
Email virus scanner, vulnerability protector, spam handler
The Horde notes/memos application
Apache module to send access logs to an SQL database
MySQL/browser-based frontend for managing Apache servers
A simple POP3/SMTP webmail CGI
A CGI script for accessing POP3 email accounts
A PHP-based PostgreSQL administration utility
A fast, lightweight, bloat-free discussion board
An HTTP server with CGI support
Provides core functionality for Web application access
A modular and portable FTP server

Sound

Audacity
AudioBurn
AudioTagTool
AutomatousMonk
BEAST_BSE
cdPlayer
DarkIce
Esound
GlobeComJukebox
IMMS
Keso
Listener

A cross-platform multitrack audio editor
A commandline audio CD burning application
An MP3 and Vorbis tag manager
Generates music from cellular automaton evolutions
A powerful music composition and audio synthesis application
Just another cd player for Linux
An IceCast/ShoutCast live streamer
The Enlightened Sound Daemon
Music jukebox with integrated CDDb aware ripping
An advanced adaptive playlist plugin for XMMS
A tracker-like user interface for the softsynth Csound
Security software that listens for and records sounds

MPFC
OpenStream
Snd
Specimen
Wavsilence
wxMusik
XMMS
XMMS-SID

A music player for the Linux console
A Web interface for MP3/ogg collections
A sound editor
A MIDI-controlled audio sampler
Split WAV files based on silence
A fast and lean MP3 and Ogg library and player
The X MultiMedia System
A SIDPlay plugin for XMMS

System

ArgusMonitoringSystem
Autosh
BitDefenderLinuxEdition
ClamAntiVirus
e2fsprogs
Foomatic
Kimono
PortableOpenSSH
Recoverdm
SCDbackup
SimpleRescueCD
Submount
Babbix

Network and system monitoring software
Monitor and automatically reestablish SSH connections
A command line antivirus scanner for Linux
An anti-virus utility for Unix
Ext2 Filesystem Utilities
Database of printer drivers and printer spoolers
An extensible and flexible network monitor
Port of OpenBSD's free SSH release to Linux
A program to help recover storage media with bad sectors
Simplified CD/DVD backup for Linux
Mid-size bootable CD (7Mb) suitable for business card CDs
A removable media handling system for Linux
An all-in-one 24x7 monitoring solution without a high cost

Essentials

Allegro
ALSA
Avifile
CheckInstall
GLib
glibc
GTK
gtkmm
Guile
Jigdo
lesstif
libESMTP
libmcrypt
Libsigc
libstdC++3
Mesa
ncurses
OggVorbis
RAWRITE
SDL
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
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
LGPL'd 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
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
Provides VGA and SVGA modes in a console

Essential disc info

Read this important information before you use your *Linux Format* coverdisc – CD or DVD. We've collated some helpful info to help you get the most from these jewels of data!

FINDING THE ESSENTIALS

MISSING SOMETHING?

As many of the programs on our discs are the very latest releases, they are often built on the very latest libraries and may depend on other packages your current Linux setup does not contain. We try to provide you with as many of these important supporting files and libraries as possible, though obviously we don't have space to include absolutely everything.

In many cases, the latest libraries and

other packages you might need will be included in the "essentials" folder on the disc, so if you are missing dependencies, this is the first place to look.

PACKAGE FORMATS

Wherever possible, we try to include as many different types of package for an installation as possible, whether that be distribution specific RPMs, debs or whatever. Please bear in mind that we can only do this where space permits and when the packages are available.

We will, apart from exceptional or legally restricted situations, include the source files for any package, so that you can build it yourself.

DOCUMENTATION

These pages provide helpful information on how to install and use some of the packages on the CD. Please note that many of the applications come with their own documentation, and there are additional notes and files in the relevant directories.

CREATING INSTALL CDS WITH CDRECORD

The quickest way to burn an ISO image to CD is with *cdrecord*. You need to be root to do this. First find the address of your CD-writer with

```
cdrecord -scanbus
```

This will show the devices connected to your system. The SCSI address of each device is the three numbers in the leftmost column, say 0,3,0. Now you can burn a CD with

```
cdrecord dev=0,3,0 -v /path/to/image.iso
```

You can simplify the command by saving some default settings in `/etc/default/cdrecord`. Add a line for each CD writer on your system (usually one) like this

```
Plextor= 0,3,0 12 16M
```

The first item is a label, after the SCSI address you put the speed and the buffer size to use. You can now replace the SCSI address in the command line with the label, but it gets even easier if you add

```
CDR_DEVICE=Plextor
```

Now you can burn an ISO image to disc with

```
cdrecord -v/path/to/image.iso
```

If you really don't want to use the command line, *gcombust* will do the job for you. Start it as root, select the "Burn" tab and the "ISO 9660 Image" gadget near the top of the window. Put the path to the image file in the gadget and press "Combust!". Now put on the kettle while the CD is created for you.

Other OS?

You don't have to use Linux to burn the ISO to a disc. All Linux-specific bits are already built into the image file. Programs like *cdrecord* simply dump it to the disk. If you don't have a CD-writer, find someone who has one, and a DVD drive, and use the CD burning software on their computer. It can be Windows, MacOS, AmigaOS whatever.

No CD burner?

What if you have no CD writer? Do you know someone else with one? You don't have to use Linux to burn the CDs, any operating system that can run a CD-writer will do the job (see above).

With some distributions it'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](#)

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 SuSE7x 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**.

ONLY AVAILABLE WITH
LINUX FORMAT
MAGAZINE

LINUX PRO

FROM THE MAKERS OF LINUX FORMAT

MAY 2004



Tim O'Reilly

Open Source, the GPL and the future of everything

PLUS

App Server 10g

Oracle's latest iteration may make you reconsider middleware

Security

Is a policy-based rather than threat-centric solution the way forward?

Storage

3ware's Scott Cleland on the future of Serial ATA RAID

Not Working?

Network management and the tools that can make it easier!

PRACTICAL LINUX SOLUTIONS FOR I.T. PROFESSIONALS

Welcome

TWENTY PAGES OF REAL-WORLD LINUX FOR IT PROFESSIONALS

Do licences really matter? When we talk about the licences generally governing 'free' (as in beer) software, as far as the end-user is concerned, the answer is generally "no". If you download some software, compile and install it, then use it for the purpose it was intended, it doesn't really matter whether it is covered by GPL, NetBSD, *Apache* or some other scheme the author made up after a heavy session in the pub. In fact, many people rarely read the licence. OK, a lot of Linux software is covered by the same licence – the GPL – which probably nearly every Linux user has glanced through once; some may even have read it thoroughly. A small group may actually understand most of it.

To developers though, it is usually of extreme importance, both in terms of the licences their own work is produced under, and the licences of any code or libraries they use in the course of developing stuff. As they are the ones creating the work, it is perhaps understandable that they are sometimes a little precious about what can and can't be done about it. To the IT departments of most companies, the licence matters too. Though some may not care to look at it, there are various licences in the 'free' software arena that may have restrictions on certain uses.

But can too many licences be such a misdirection, a burden on the time and care of users and developers alike, that we end up focusing too much on the wrong issues? There are some XFree86 developers who think so, and our special interviewee this issue, Tim O'Reilly, agrees. Find out what he has to say on the GPL and more diverse topics in our exclusive interview, starting page 10.

Nick Veitch Editor
nick.veitch@futurenet.co.uk



A lot of Linux software is covered by the same licence – the GPL – which most Linux users have probably glanced through once. Some may even have even understood it...

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 matthewnailon@futurenet.co.uk

Publisher

Kelley Corten

Photography

IBM, Thinkstock, Pixtal

ADVERTISING & MARKETING

Senior Sales Executive

Andrew Hudson

andrewhudson@futurenet.co.uk

(01225) 442244 Ext 2441

Copyright No part of this publication may be reproduced without written permission from our publisher. All letters & contributions to *Linux Format* are submitted and accepted on the basis of non-exclusive worldwide licence to publish or license others to do so unless otherwise agreed in advance in writing. *Linux Format* recognises all copyrights in this issue. Where possible, we have acknowledged the copyright holder. Contact us if we haven't credited your copyright and we will always correct any oversight. We cannot be held responsible for mistakes or misprints.

LINUX is a trademark of Linus Torvalds, GNU/Linux is abbreviated to Linux throughout for brevity. All other trademarks are the property of their respective owners.

LinuxPro is a part of



magazine

Published by



Media with passion

30 Monmouth Street, Bath BA1 2BW
Tel +44(0) 1225 442244

Front cover: **Tim O'Reilly**, courtesy IBM
(copyright IBM)

CONTENTS

STORAGE

3ware and the SATA controllers of the future

p4

SECURITY

Having a policy can save a lot of effort and confusion.

p6

NETWORKING

Give yours a health-check

p8

TIM O'REILLY

Open Source and the GPL

p10

APP SERVER 10G

Oracle have more to offer than just databases.

p19



04

10

STORAGE: SATA

3WARE AND SATA RAID

Serial ATA is the storage medium of the moment. **NICK VEITCH** talks to leading enablers 3ware about the future of this technology.

Serial ATA hasn't exactly taken the world by storm, but it is definitely a strong technology with a future. Shrugging off the image of ATA being the shabby second-cousin of SCSI may be a bit of a struggle to begin with, but with robust features and the commitment of companies like Seagate, Western Digital, and of course 3ware, it seems that IT businesses are opening up to the possibilities it offers.

3ware has been involved in the Linux community for a long time, and is most notable for its support of RAID controllers for ATA drives. Last year, the Escalade series of controllers took the company into SATA RAID, and with new products lined up for 2004, it seemed like a good idea to catch up with **Scott Cleland**, 3ware's director of marketing.

LINUX PRO: You now provide a great many solutions for SATA – is that something you see as complementing existing RAID tools, is it a new market, or is everyone going to migrate to SATA RAID?

SCOTT CLELAND: Certainly for us, SATA is our core business. We were the first to market with a high capacity SATA RAID controller in 2002. We've gained a lot of experience in the high port ATA RAID environment. As we've moved up the food chain in the enterprise, Linux has been a key enabler for us to get into these markets. I'm not beating about the bush – 3ware went to market with Linux as the core OS supported, rather than going after the traditional storage plays in the Windows environment. 3ware saw an opportunity to leverage the demand for high-capacity storage, and Linux was a big part of that, because of its popularity in those markets, like High Performance Computing.

3ware's global presence in the storage market is reflected by the range of international options on its web portal.



We gained really good momentum around large installations like CERN, and a lot of labs in the US. We made a name for ourselves in that space. Linux turned up everywhere, so we gained momentum in environments like NAS too. As SATA became popular as a complimentary technology to enterprise SCSI and fibre, we became more involved with disk to disk backup. We have a lot of traction in these vertical markets where Linux is popular. And of course Linux is now meeting Windows in the more traditional enterprise markets.

LXP: There are many reasons for using SATA, but there are still reasons why people opt for other solutions. The capacity of current SATA drive for example isn't great. Do you think that is impeding the take-up of SATA?

SC: I think as the IT community becomes more comfortable with SATA, we'll see more deployments. In some ways, SATA has the (false) perception among many IT decision-makers of being a desktop drive. We leave the deployment in the hands of our customers. We try to offer the highest level of redundancy and reliability with our RAID controllers, and that's a very important factor. 3ware has taken a great deal of time in engineering techniques and methods to make sure that your data is safe – we use a SCSI model, so we have features like dynamic sector repair, disk-scrubbing, you can protect sets where you have a spurious POR (power on reset) to a drive with data still in the cache. We do a lot to make sure you have a robust solution around the drives.

LXP: In many ways, ATA has connotations of being 'the poor-man's storage format'. I understand from things we have researched before that the will and intention of those participating in SATA is to incorporate the very best features possible in SATA.

SC: SATA can't use the same mechanical parts that in SCSI and fibre channel drives do. You can't burn-in the drives as you would with SCSI. A burn-in on ATA is nine hours – so you're between a rock and a hard place: you want to add robust features, while keeping the cost of the ATA space.

It's an interesting challenge, and I think 3ware has made a name for itself as being up to the task to solve those problems – offering as many of those reliability features on the controller and outside of the drive.

LXP: Do many people use mixed storage solutions?

SC: Sure. There are many different ways of doing it. I guess

the principal environment that is used is in tape backup. SATA is used as a near-line augment to the tape backup, so you can offload the data to SATA and backup from there. There are also high-end systems using SATA and fibre channel in the same environment, with the fibre being used for database access and SATA for near-term storage.

LXP: You mentioned that Linux is key to your business model. The problem is: there are still a great many incompatible versions of Linux: does that cause any issues? We notice you support Red Hat and SUSE...

SC: It's definitely a challenge. It's tough to run Debian for example in our QA department, because it's often changing. The commercial distros are an easier model to test, and SUSE and Red Hat have the lion's share, though in 2004 we will be testing on Miracle Linux, Mandrake; then we'll have to get on the Fedora bandwagon.

LXP: Debian is interesting because it seems to crop up quite a lot in the server space. It seems well regarded because of the robust package management, but I can see that it may be a problem...

SC: Fortunately we have been in the kernel for quite a while. I know we haven't lost business because we are not supporting a particular flavour. For instance, in the 64-bit space, we have had a driver in the kernel. We haven't produced a 'package' for 64-bit yet, but we do support it in the field. We have a product being released in April that will be tested officially on AMD64 and released through QA. Typically RAID companies will have one or two Linux gurus, and we have a particularly good one who understands how the community operates and is very active on the forums. Plus, since our architecture was designed from the ground-up with Linux in mind, it works very well for us. Other RAID companies have difficulties when they move to producing Linux drivers, as their IP will be tied up in the kernel driver.

LXP: The opinion that I hear coming back from the community, is that 3ware is seen as very much part of that community, rather than a company trying to use the community as another way of selling hardware.

SC: It was a good decision that 3ware made to embrace Linux. We don't track it percentage-wise, but it's probably safe to say that more than half of our business is in Linux. We are also doing things like looking at a Linux BIOS – we're trying to stay on the bleeding edge of where we need to be.

LXF: Since Linux continues to be the fastest growing OS, you must be confident of an expanding business.

SC: Oh yes, were working with a number of businesses. We see that in terms of capacity, our largest opportunity is with the supercomputing community. And there are plenty of other vertical markets that we are focussing on: NAS, disk to disk, a whole load of stuff.

As you know, we were acquired recently by AMCC, which brings to us a whole new level of opportunity in terms of a breadth of engineering: an ASIC company has lots of tools

we can leverage. We'll be combined with a fibre-channel group. The opportunity to provide products in the NAS and SAN space becomes real for us. We're also working towards building iSCSI solutions for the Linux space.

LXP: There now seems to be a bewildering array of storage options. Is this proliferation a good thing, or will people eventually decide that one or two solutions best suit the majority of needs and the rest will fade?

SC: That's difficult question. The reason we go after vertical markets is that that's where the growth is. We will continue to scale our business in these verticals. I don't see a situation with one or two suppliers. There will be consolidation, but a lot of solutions haven't yet been fully developed.

We have a lot of partnering to do with companies doing virtualisation, provisioning, interesting iSCSI markets. I think in the 1Gb space the iSCSI solutions will develop into an alternative for direct attached storage DAS. We're leaving the SAN for the fibrechannel boys for the moment.

LXP: The HPC market has really taken off. So not only is that a strong segment, but it's growing...

SC: The best thing about the HPC market is the economy: prospective users are looking to bring in the lowest cost-to-performance ratio they can. From a storage perspective, they want something as low cost as they can get with the best spec they can get, and that's where we've had a lot of interest. They may have a requirement in 2004 for a pipe sustaining 300MB/second, and we can offer that right now. That's why SCSI and fibre channel are losing some traction.

LXP: At the moment you offer SATA controllers ranging from 2 to 12 ports. Just out of interest, whereabouts in that range is most of the business at the moment?

SC: It's rapidly moving towards 8 ports. We attach approximately 50,000 SATA drives a month. When the transition started between parallel and serial ATA, we moved from a 4-port sweet spot to a 5.5 port spot. It's going in the right direction!



The popular Escalade series of controllers may soon make way for a completely new hardware platform.

“MORE THAN HALF OUR BUSINESS IS IN LINUX. WE’RE ALSO DOING THINGS LIKE LOOKING AT LINUX BIOS, WE’RE TRYING TO STAY ON THE BLEEDING-EDGE...”

LXP: Are we going to see any larger devices?

SC: Well, we have a new product coming out soon that we hope will enhance our position in the HPC market. The great thing about 3ware is that we've developed ATA RAID from the ground up, either as an augmented or complimentary or alternative solution to traditional storage. Other people have delivered ATA RAID because they had to. Their featuresets, while seemingly broad, may not be the features that customers require in the ATA space, and it's then very difficult to effectively port a feature set from SCSI to ATA and have it work the same. We have learned a lot from being only in the ATA space. ■■■

SECURITY

MANAGE YOUR NETWORK VULNERABILITIES THROUGH POLICY-BASED SECURITY

Solsoft CEO **GILLES SAMOUN** shows how security and communications are inextricably linked with *your* bottom-line.

In today's digital economy, if your network is down, your business is down. To protect corporate networks, IT managers in 2003 devoted an increasingly large proportion of their IT budgets to IT Security. Even so, there were almost as many Internet security incidents in the first six months of 2003 as there were in all of 2002, according to the CERT Coordination Centre.

If security spending is up, yet security breaches are on the rise, there is clearly something wrong with the way we allocate our security dollars. The old approach of adding new layers of security to meet new threats – such as adding anti-DoS (denial of service) devices, virus protection programs, and intrusion detection systems – is clearly not adequate since these new layers of security have a limited scope. While a device may prevent a class of DoS attacks, it cannot preempt or predict new attacks, and since threats are constantly evolving, this is a stopgap measure at best.

How, then, do we go about correcting this situation? First, we must change the way we think about security and communications, linking the two, rather than treating one as mission-critical and the other as an afterthought.

REACTING TO VULNERABILITIES

Consider the recent DoS vulnerability that was discovered on Cisco switches and routers running IOS. The vulnerability was announced by Cisco, and shortly after relayed by CERT on July 16, 2003. It caused unprotected Cisco devices running IOS to stop processing inbound packets on ports 53, 55, 77, and 103. Two days after the Cisco announcement, on July 18th, an exploit was published on an open mailing list, providing instructions to malicious users on how to easily spread the packet sequence that would create a DoS attack and disrupt the operation of unprotected IOS devices.

Several businesses were impacted by this vulnerability. Of the reported incidents, some were due to direct attacks, while others were caused by network managers trying to fix the problem. The hasty patching of hundreds of devices at a time created more problems than the original vulnerability itself!

In today's larger enterprise, a distributed, often multi-vendor environment creates a management nightmare. Each device must be reconfigured manually; typically, different types of devices – even those from the same vendor – have different management interfaces. In globally distributed networks, by the time a corporation has learned about an exploit, crackers

will have a huge window of opportunity before that company's many devices are reconfigured. And given that reconfiguration is not only a labour-intensive process, but also a tedious and error-prone one, the determined cracker can be fairly sure that the exploit will continue to exist even after patches are applied, albeit on a much smaller scale.

In order for a network to keep up with new vulnerabilities and threats, a network administrator needs a management solution that allows them to address the many security devices on the network directly, to push patches and updates in an automated, uniform fashion, while also receiving verification that the new configurations have indeed been executed.

A new class of software – security policy management software – does just this. Unlike point-to-point management, where security devices are configured one-by-one across the network, policy-based management allows network administrators to closely follow business practices and requirements. Sound security demands global consistency. For example, if you modify a rule on a single device without understanding the global policy implications, this action could have a negative impact on business application flows and even compromise security of the entire network.

A POLICY-DRIVEN APPROACH

In the case of the Cisco vulnerability, a policy-based system would have headed off the vulnerability well before it was ever a target for exploits. If your enterprise is running switches and routers with ports 53, 55, 77, and 103 open, then your security policy was faulty. Why? Because those ports shouldn't have been open to exploit in the first place.

Having these ports open is no different than the person who leaves his keys in the ignition of his car, yet is surprised when the car is stolen. Ports 53, 55, 77 and 103 are ports that very rarely (if ever) need to be used, and therefore have no reason to be opened. The main reason they are open on a vast majority of devices today, is that most businesses follow the conventional security model that values communications over security.

The conventional security model says that all traffic that is not explicitly denied is automatically authorised. But why does the conventional model function in this insecure manner? Because it is so very difficult to manually configure the many disparate devices on the network. This leads

network security professionals to leave everything open by default, only denying access to ports that are notoriously vulnerable and listed as such in security books and device configurations manuals. Until July 16, 2003, ports 53, 55, and 57 were not part of that list. As long as companies use this approach to security, they will be vulnerable to new attacks, especially those companies with large, distributed, heterogeneous networks. On the other hand, companies with security policy management software in place will be alerted to these vulnerabilities. They will have an automated means for correcting them, while also receiving automatic reports that verify that the corrections have occurred.

A point-by-point approach to defining security leaves the network vulnerable through back doors, misconfigurations, and even branch offices and telecommuters. A policy-driven approach defines what exactly can pass into and back out of the network, no matter where that traffic is coming from or what its destination is.

Instead of leaving ports open by default, a security-conscious company should instead assume that by default everything that is not explicitly authorised is automatically denied. Every authorisation that is defined is therefore a calculated risk that leaves little space to the unknown. This approach would protect companies from new threats, such as the IOS vulnerability, the SQL Slammer, and other vulnerabilities that use exotic ports that have no business being used in the first place.

Critics of this approach will argue that this creates a burden on network administrators, forcing them to define many different types of permissible communications. This may seem like a daunting task, especially when different business units within the company may well have different requirements. However, when compared to the task of recovering from an attack, the burden of defining permissible types and classes of traffic is mitigated – at the very least, this process does not bring your business to a grinding halt, the way an attack might, and, in the long run, it could very well keep your business up and running when new attacks emerge.

Often, knowledge of a problem is considered half the battle, and that would be the case with policy-driven security if companies were building new networks from the ground up. Instead, they are trying to secure existing networks based on hundreds or thousands of devices, and they need to be able to integrate new devices as the network evolves and grows. Therefore, in real-world deployments, the policy-driven approach must be integrated into the network itself.

SECURITY POLICY MANAGEMENT AS A NETWORK LAYER

Since each class of device, from router to switch to firewall, requires a different management interface, I propose addressing this problem through a management abstraction layer that resides above these many devices. This new software will take the configuration and policy burden away from already overworked network and security administrators and apply it to each and every network device in a logical and consistent fashion. This security policy management software helps administrators create a global security policy; it discovers what devices are on the network; and it distributes policy automatically.

Having this virtual layer means that changes to complex networks can be done in a matter of minutes – whether one or hundreds of devices need to be reconfigured – enabling rapid response to network events or new business requirements. IT personnel can build and schedule security scenarios to meet business needs, applying policy changes on the fly and simplifying what would normally be a complex, manual, and error-laden process. Typical examples are after-work policies that are made restrictive and work-day policies that are left more open.

Security policy management software is just now starting to hit the market. While the benefits may seem revolutionary to the companies deploying it, this software is actually evolutionary, a logical descendant of business process applications such as CRM (Customer Relationship Management) and SFA (Sales Force Automation). After all, what do these applications do but establish policies and make otherwise chaotic systems manageable?

POLICY-BASED SYSTEMS CAN PROVIDE SYSADMINS WITH A UNIFIED MANAGEMENT INTERFACE FOR ALL NETWORK SECURITY DEVICES

Companies in the security policy management space are seeking to give security administrators the same level of process control that their counterparts in sales or customer service already have. But whereas SFA benefits only the company and its bottom line, security policy management software has broader implications: the more secure the basic constituents of the network are, the more secure the network as a whole is, including all of its applications.

With security policy management software in place, network managers now have visibility into and control over their networks. If a Fortune 500 company realises that Instant Messaging is posing a security risk, that company now has the ability to disallow that form of communication. With traditional networks, shutting down IM traffic is a nearly impossible task. How do you execute this order, and once the order is sent, how do you verify compliance? In a policy-based system, this rule is readily available and the implementation is seamless. Security policy management software governs this process, from the generation of the policy through the execution of the change in configuration to the reporting that verifies that the policy is in place – and the process is the same whether you have one device or hundreds of them from different brands. Network administrators no longer simply react to problems. Instead, the policy-based system provides them with a unified management interface for all security devices. It enables network administrators to be proactive, ensuring that best-practice security rules are enforced without human intervention on each device across the global enterprise.

Networks are constantly changing to meet evolving business needs, and only a policy-based, process-driven approach to security will ensure that changing networks remain secure networks. Anything less puts your company, your customers, and your business partners at risk. ■■■

AUTHOR INFO

Gilles Samoun is CEO of Solsoft, a developer of security policy management software that serves as a centralised platform to automate the provisioning of network security rules on multi-vendor routers, switches, firewalls, and VPNs.

Solsoft is exhibiting at Infosecurity Europe 2004. The event brings together professionals interested in IT Security from around the globe, with suppliers of security hardware, software and consultancy services. Now in its ninth year, the show features Europe's most comprehensive FREE education programme, and over 300 exhibitors at the Grand Hall at Olympia from 27 to 29 April 2004. www.infosec.co.uk

NETWORK NOT WORKING?

NICK VEITCH reveals some handy tips on network management.

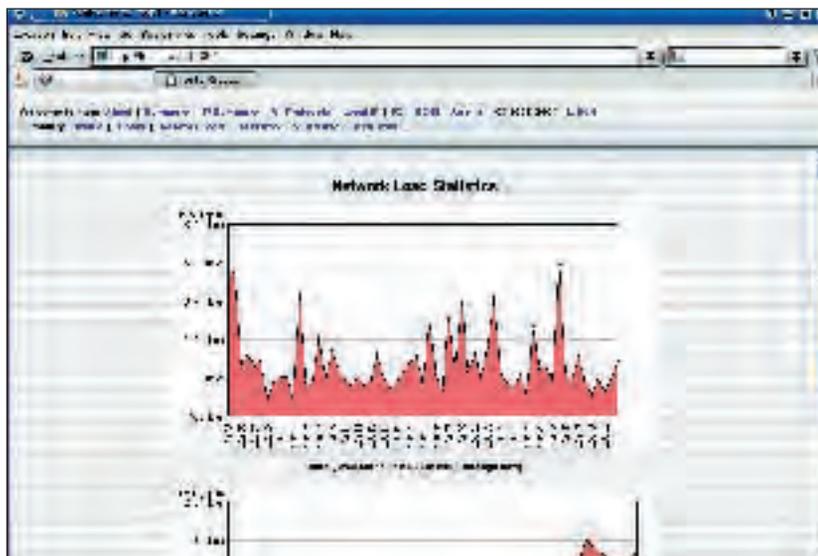
As with every other IT resource, network bandwidth isn't free, and neither is the server power required to deal with the traffic. And yet, it's amazing the number of admins who will simply add more of either (or indeed, more of both) to their network solution without ever wondering whether all the traffic that causes the demand is either necessary or desirable. When your corporate network is actually being brought to its knees by employees using instant messaging software, you'd be better of looking at an AUP rather than more bandwidth. Or the problem may be more sinister – a compromised server running a DoS attack, or a misconfigured mailserver acting as an open relay for spammers.

The plain truth is, you don't actually know what is happening on your network unless you look. All those

packets flying backwards and forwards can tell you a lot about network utilisation, and even help you understand the way your organisation uses other resources.

IT'S ALL VERY WELL TO PING A REMOTE FILESERVER TO GET TTL & ROUNDTRIP RESPONSES, BUT THAT DOESN'T MEASURE EFFECTIVE THROUGHPUT

Understanding load statistics can be key to management.



NTOP

ntop is meant to be the equivalent of the well-known monitoring tool, *top*, but written especially for networks. Making use of *libpcap*, *ntop* runs on any machine on the network and captures network packets. The packets themselves are of no interest, but all the header information is kept and stored in a database. From this information, *ntop* can then analyse the data and produce useful information on network traffic.

Although designed to be used primarily on a network-attached computer, it is also possible to capture data from other sources – eg some routers and switches which allow packets to be captured. This can then be 'fed in' to *ntop* to provide a clearer picture of the traffic through the router. To analyse data produced using *tcpdump*, simply supply the `-f` switch when running *ntop* on the command-line, followed by the dump file.

Conversely, from the web interface it is possible to dump out the data logged so far by *ntop* in a variety of formats including XML, text, and even Perl, PHP and Python formats, the latter three of which produce syntactically

correct array structures of the relevant data, for possibly use in other scripts.

The application uses its own built-in server to provide remote access to the collected data. Even on a local machine, for most tasks it is more convenient to run a browser and connect to the server (which you'll find on port 3000 by default).

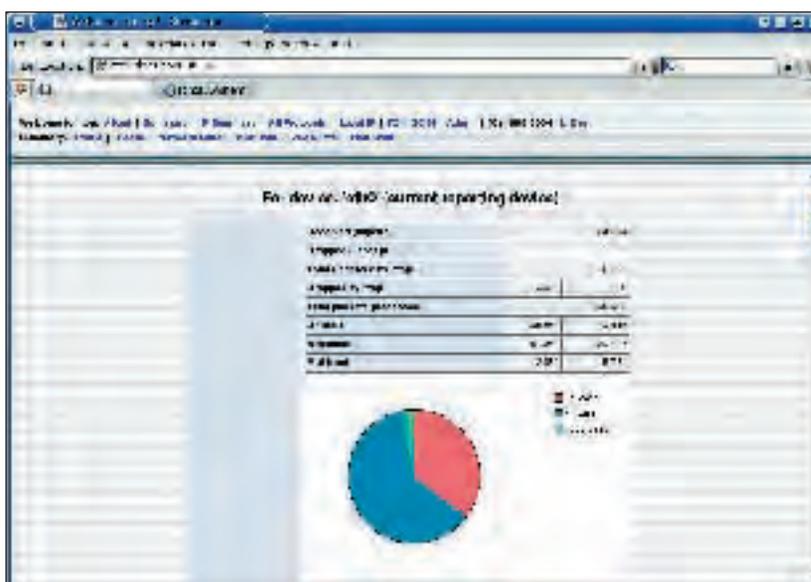
LOCATION, LOCATION

One of the crucial things to consider for *ntop* to be effective is choosing its exact position on the network where you run it. For some purposes – for example, identifying problems with the amount of broadcast packets harassing the network – obviously anywhere within a segment will do. To effectively determine problems and bottlenecks with file servers, gateways and routers though, you really do need to analyse the captured data from the relevant devices.

This is where a proper network analysis tool comes into its own. It's all very well being able to ping a remote fileserver and establish what sort of ttl and roundtrip responses you get, but that doesn't measure the effective throughput. While problems on a fileserver are more likely to be down to issues with the server itself – inefficient storage solutions being top of the list – it helps to be able to prove that the speed and general connection of the network is not impeding performance. *ntop* can provide a monitor of the effective throughput through the life of a large file transfer for example, which on the one hand will give a more realistic measure of the bandwidth consumed/available, it will also be able to highlight any problems. As an example, running *ntop* on the client end and looking up the Host info page will list – amongst other details – the total data sent and received, including the number of retransmitted packets.

SECURITY

ntop can also act in the role of a primitive NIDS (Network Intrusion Detection system). Because it knows where packets are destined, it can quite easily spot traffic headed to 'bad' ports, and flag up unwanted requests (such as DoS packets) either travelling into – or going out of the network. Obviously many other tools exist for specific security



ntop monitors packets on a device level basis, and includes more statistics than you could ever reasonably need.

monitoring of network packets, for example *SNORT*, but this is a useful extra bonus for *ntop* users.

Iperf

Running *ntop* to locate and analyse potential bottlenecks is useful. Often though, it is handy to analyse a specific end-to-end connection speed. By this we don't mean ping, which, as discussed earlier, really only tells you that a host is accessible, rather than anything particularly meaningful about the connection to it.

Iperf is a fairly simple tool that can perform a more meaningful and repeatable test than trying to ftp between two computers and timing the results. Although in some ways analogous, *Iperf* – while remaining simple – has more than enough options to successfully test end-to-end connections in a thorough fashion.

Iperf works as a server/client pair – an *Iperf* server must be installed on the endpoint you wish to test, with the client at the other end. The server will listen out for connections (on the default port of 5001) and results of the transfer will be displayed periodically on the client end.

The time and frame size of the data can be changed to simulate different types of data transfer, which may help highlight potential problems. Usefully, it is also possible to test different TOS packets. TOS (Type Of Service) is used by routers to determine the priority for handling packets, and different types can be used to request particular performance – reliability, throughput or consistency for example. As *Iperf* can also be used to test UDP frames, this combination allows to test the reliability of VoIP or video streams, or indeed better assess what sustainable bandwidth is available for these services.

MORE TO DO

Of course, this only scratches the surface of network monitoring and the huge range of tools available. The important thing to realise is that the tools do exist. We may have had *tcpdump* for years, but there is a difference between merely collecting the data and making sense out of it. The health of your network is important, and you could manage it better with a little help! ■■■

Iperf LIBRARIES

Another useful add-on with *Iperf* is that the library used by the software can be quite useful for coding your own applications – for testing purposes or otherwise. This is reasonably well documented and could form the basis of more customised network testing software.

INSTALLING AND USING NTOP

The recommended way to get hold of *ntop* is to download the source through CVS. Try the following:

```
export CVSROOT=:pserver:anonymous@cvs.ntop.org:2401/export/home/ntop
```

```
cvs login
```

```
cvs checkout ntop
```

Enter **ntop** as the password if asked. Then descend into the *ntop* directory and then do the usual:

```
./configure
```

```
make
```

make install

There are also a number of packages available to download from the website (www.ntop.org/ntop.html), including some binaries for Windows, but there is no guarantee these will be the latest version.

Before you install, *ntop* also requires *libpcap* and *gdbm*. Both of these are usually present in most distros. The configure script will in any case let you know if you have an incompatible version and point you in the right direction to find one.

COVER FEATURE **TIM O'REILLY**

INTERVIEW

THE O'REILLY MONOLOGUES

We invited **TIM O'REILLY** to talk to *Linux Pro* about Linux: instead he shared his thoughts on Linux, Open Source, the GPL, Microsoft, and more. By his own admission, "It's a lunchtime rant"...

TIM O'REILLY: To me, the biggest story with Linux – and the subject of the talks I have been giving over the last year – is the way that the Internet is changing the whole computing paradigm. A lot of the Linux advocates and a lot of the Free Software thinkers are backwards-facing! I have this trick question I use in my talks: “How many people in the audience use Linux?” – depending on the audience, it'll be 20 or 80 per cent. But when I ask, “How many of you use Google?” – every hand goes up. And my rejoinder is: “In that case, you all use Linux.” We have this backwards-facing idea that what you use is limited to what sits on a desk in front of you, and that's just not true any more.

We're moving into a new era. There's a wonderful quote that I love from Dave Stutz (author of *Rotor*, the shared-source implementation of .NET). In his open letter, written on leaving Microsoft, he ends by saying, “software written above the level of a single device will command high margins for years to come – stop looking over your shoulder, and invent something!” His letter was all about how Microsoft is fixated on Open Source licenses, but that MS needs to get used to the fact that Open Source is commoditising a lot of the areas that used to be valuable in software. Meanwhile, there are vast new fortunes being created by people building on top of Linux – Amazon is built on Linux, Google is built on Linux, Yahoo is built on FreeBSD, www.SalesForce.com is built on Linux. What does that tell us? Why is the Linux market not owning its greatest successes?

Google: Linux's killer app If you talk to most Linux advocates, they'll say, “Isn't GNOME going to be great when it's as good as Windows?” To which I reply: “Why does that matter? It doesn't matter that much.” The future is in software that is living on this new platform that we're building on the Internet, and an awful lot of that platform infrastructure is Linux already. I don't see a lot of focus in the Linux community saying, “Let's go do a story on how Google has the world's largest computer installation, all running Linux.” Google has tens of thousands of Linux machines to facilitate its great application – the killer app of today's computing society. It puzzles me, and it frustrates me that the Linux community isn't saying, “Never mind Red Hat – look at our successes. Look at Google; look at Amazon.”

Part of the problem is the political focus of a lot of people who speak for Linux. They go, “OK, these guys [Google, Amazon et al] aren't releasing their software as Open Source, therefore they're not part of our community.” I would love to see the Linux community focus more on these killer apps. www.SalesForce.com is a good recent example. It isn't a Google or an Amazon, but it is about to go public with a billion-dollar IPO; yet why is no one trumpeting the fact that it's running on Linux?

In short, I think it's really important to recognise that people are building these next-generation applications on top of Linux, but they're not Open Source, they're not constrained by any licences because they're not distributed software. I'll come back to that thought in a minute.

Software commoditisation The rules are changing. And if you look at a computer industry history, you can understand a bit about what the new rules are going to be. The last



One day, will all businesses will be this aware of their place in the community and the consequent ethical, environmental and social responsibilities?

time this happened was with the personal computer: back in 1981, IBM changed the rules of the computer industry. Up to that time, all computer hardware was closed and proprietary. The PC was the first open hardware architecture. IBM thought it knew which particular rules were changing, but the rules then changed way more than IBM thought they would. IBM thought that what it was doing was, “OK, we're going to create this commodity personal computer, and it'll be sort of like VHS versus Betamax, and we'll own the market.” But what IBM didn't realise was that first of all, everybody could make the same commodity hardware, and consequently the rest of the market would eventually be bigger than just Big Blue. IBM also turned over the software to Microsoft, because it didn't realise at the time that once you commoditise the hardware, software becomes way more important.

“IT'S REALLY IMPORTANT THAT PEOPLE ARE BUILDING THESE NEXT-GENERATION APPS ON TOP OF LINUX; BUT THEY'RE NOT OPEN SOURCE, THEY'RE NOT CONSTRAINED BY ANY LICENCES BECAUSE THEY'RE NOT DISTRIBUTED SOFTWARE”

I think one of the consequences of Linux and the Internet is that software is becoming commoditised – open standards lead to commodity software. Whether it's Open Source or closed-source, you can't necessarily make money in the same way. Taking *Internet Explorer* as



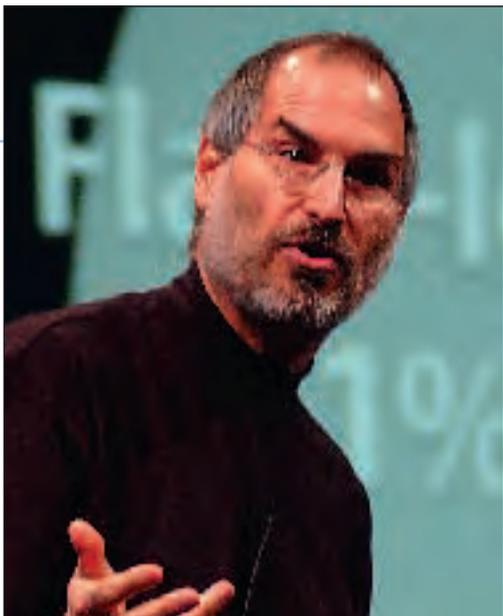
COVER FEATURE **TIM O'REILLY**

“An example – it's got to be free, because it's basically implementing an Open standard. If it's not free someone else will provide the same functionality for free. It's not really complex software – it's the system that's complex, and the system is Open. So what we have done is created a situation where the software is commodity software, just like how the PC was commodity hardware. So where does the value go?”

The value goes up the stack to data. If you look at the big successes of the Internet era, they are companies who have built these data-rich services – eBay, Amazon, Google, and now www.SalesForce.com. All these companies are building this infrastructure on top of Linux – among other Open Source tools and languages – but neither they nor the typical Open Source advocate think of them as being a part of the Open Source community.

Giving back to the community Part of what I have been trying to do with these companies is educating them about their debt to the Open Source community, and trying to help them think through how to give something back so that they keep the virtuous circle going. When Jeff Bezos (founder and CEO of Amazon) did the One Click patent thing, I said to him, *“Look, you got all this benefit from Open Source and Open standards, and now you're saying ‘the party stops here.’ Well, that's a bad idea for you, because eventually you won't have the kinds of innovations you've benefited from, and you're going to end up beholden to a commercial supplier – if you were running on Microsoft, they'd be squeezing you by the balls right now. Instead you have this Open platform, so keep that party going – figure out how to give back; figure out how to have some level of engagement.”*

I've been lobbying Amazon the most, because I know the people there pretty well, but it's not easy to figure out – even with a lot of goodwill. How does Amazon give back? Well, it uses Perl and Mason, and Amazon has started being an active contributor to the Mason community. That's a start, but Amazon is also very proprietary about its core competencies; even if it wasn't, Amazon isn't like a typical software project where you can just give away the code and



Apple's Steve Jobs: licensed the One Click patent from Amazon – something that O'Reilly has thought about doing too.

Eric S Raymond: “know that the Nutshell Guides are but the outermost Portal of the True Enlightenment.”



you have it up and running. So now, Amazon has started thinking about issues like: *“How do we open up?”*

What that question led to was the Amazon web services API, which allows some reuse. You can create interesting new applications that reuse the API's functionality. No, it's not Open Source, but in a lot of ways, it's *better* than having the Amazon code – because if you had the Amazon code, you still wouldn't have Amazon. Pieces of it would be useful, but when a paradigm changes, the rules of business really do become different. So much of the Linux and FOSS advocacy communities are still playing by the PC industry rules, when really what we should be doing is building a new set of rules that is totally based around the next generation of software.

LXP: Has Amazon actually relinquished that patent?

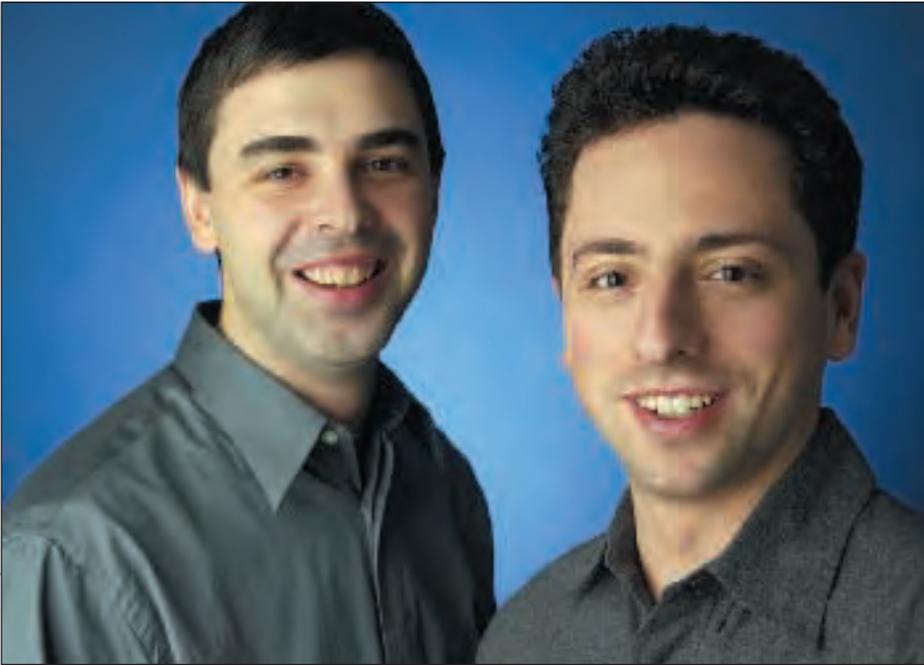
TO: No, it hasn't. Amazon did settle the case with Barnes & Noble, and I don't think it will actually use it offensively again. I think that's probably good enough... I mean, everyone's patenting; I don't think that's going to go away as long as we live, but the question is: do people actually sue other people over the patents they hold? I don't think they'll do that again.

LXP: Apple licensed it, didn't it?

TO: Yes, Apple did, and I think it should be licensed to more people. I've actually been tempted to do it myself, because I think it's a cool piece of technology.

Returning to the broader discussion, I'd like to think about how to take the ideals of Linux and the Internet – which are about collaboration, about openness, low barriers to entry, extensibility – and separate them from the original Open Source definition, which is so heavily licence-focused. If you look at the values in Open Source, there are four or five different sliders, and you can move them differently.

A good example there is the company that Brian Behlendorf and I started called CollabNet, which your readers may know about. CollabNet started out doing Open Source projects, and it still provides infrastructure for *OpenOffice.org* and a bunch of other Open Source projects,



but its biggest customer is HP, who use all the same technologies that *OpenOffice.org* uses (and *Apache* uses), but HP uses it inside its printer division – there's 3000 HP engineers collaborating inside the firewall. Completely independent of the licence axis, they are using the collaborative tools of the Open Source community.

LXP: Is that a bad thing?

TO: I don't think so – I think it's good because it's one of the real benefits. You *can* separate the idea of what licence do you release software under from how you develop it.

Looking at that axis of collaboration rather than licensing is one of the things I'd like to get people to think about. And it's more than just collaborative development tools and processes: look at the whole architectural aspect of Open Source – you design a system in such a way that it is easy for people to work independently. I originally got that idea from a conversation that one of our editors had with Linus Torvalds back when we were putting together the book *Open Sources*. Linus said something along the lines of, "I couldn't have done what I did even if I had the Windows source code because Windows isn't architected in such a way that I could have taken off a piece and worked on it."

When you think about it, Unix has this 'small pieces, loosely joined' architecture, which meant that you could bite off a piece that's the size of a small team and replace it, while someone else could work on something else completely different. And that's a whole architectural design that just isn't present in a system like Microsoft Windows.

I had a similar conversation with someone at Apple about some of the things I wished that company would do. For instance, I love *Rendezvous* – a very cool Open Source technology that came out of Apple. But never mind the Open Source part – full advantage of *Rendezvous* isn't even being taken within the Apple applications – and that's just the tip of the iceberg. Apple has got some wonderful things in its various iLife applications: there are buddy lists in the Address Book, and also in iChat; and to return to the *Rendezvous* example, it

Google founders Larry Page and Sergey Brin: probably the best-known Linux advocates in the world...

was originally in some of the apps but not in others. Why doesn't Apple just let end-users get all this stuff across all the apps and make it consistent? They all have great functionality, and I'd love to be able to say, "these are the people on my buddy list with whom I want to be able to share my addresses"... The reason why this doesn't happen already is probably because the developers hadn't thought of it. Andy Hertzfeld, when he's playing with *Chandler* (a networked PIM) he's thinking that way. He's thinking, "Network!" How do you rethink these things in the age of the network? You think, "Oh yeah, of course I'm connected to these other people".

I'm totally rambling here, but when you look at all the social software stuff, you know – like Friendster or Orkut. They are all hacks to get around the fact that we can't share our addresses very easily. Apple says: "We didn't architect the applications that way," and: "We have to totally refactor these applications to make them into services that can be used by other apps". That's what I mean by an architecture that supports re-use.

Licensing and Unix heritage Linux has that Unix heritage where everything is designed to be a service to something else, and I don't think that people give that enough credit in understanding what's important about Open Source. That's way more important than licensing. For instance Unix, under the old AT&T licence, had a community that looked like Open Source, until AT&T did something stupid, which was to shut it down. So to me, licensing is just housekeeping you do after the fact to say, "Don't be stupid!" I actually believe, long term, that none of the Open Source licences matter that much, because what will happen is that people will realise that all these things that the licences try to enforce are really good practice.

We won't need the licences because companies will realise that that's the right thing – the smart thing – to do. As we move from the religion of Linux to the *science* of Linux, we'll get to a point where everybody says, "Oh...". If

"IF YOU WERE RUNNING ON MICROSOFT, THEY'D BE SQUEEZING YOU BY THE BALLS... FIGURE OUT HOW TO GIVE BACK... TO OPEN SOURCE"

you're a pragmatist like me, you'll see that sometimes people do in fact sometimes try to keep things proprietary, one way or the other. Contrary to the point of view held by Richard Stallman, I don't believe that everything needs to be Open, that everything needs to be Free.

I'm looking to maximise value, and I see that everyone from Red Hat and CollabNet, to people like Google and Amazon, is making choices about where it's going to be open and where they are going to keep things to themselves. I make those choices myself – sometimes we make books free, or we make parts of books free, but if we made other things free we'd be a smaller business – or we might even be out of business, and then the value that we are able to create wouldn't be happening. 

COVER FEATURE **TIM O'REILLY**

◀ **Maximising value** Larry Wall has this great quote that's relevant here: *"Information doesn't want to be free; information wants to be valuable."* And for me, the real question is, *"How do you maximise value?"* I believe that one valuable thing Open Source teaches us is that people were not maximising value in a whole bunch of ways. Companies would create something useful that didn't succeed in the marketplace, and they'd lock it away somewhere, kind of like the end of *The Raiders of the Lost Ark* – in the big storage vault where no one will ever see it again. That's not maximising value – giving it away is maximising value because some dividend *always* comes back to you.

Larry Wall creates Perl, or Linus creates the kernel. If he goes off and tries to start a company, it's going to divert his energy: he's probably not going to succeed at raising money, he's not going to get to do what he does, so instead he maximises value by saying, *"Hey, I've got something useful and I'm giving it to you, if you do something useful, give it back."* We have this wonderful gift culture that grew up, and it works, right? But it's just one thing that works.

We also know that Bill Gates is creating a lot of value. Anyone in the Linux community who denies that is stupid, to be quite honest. I mean, Bill Gates has created huge monetary value for his shareholders, and he has created huge value for himself.

He also created the ecosystem that Linux depends on; I don't think Linux would be possible without Microsoft, because Intel and Microsoft created the monoculture of machines that were all kind of saying, *"Hey, we have this commodity hardware out there, let's now put some commodity software on it."* It's a natural evolution. What Bill Gates did wrong was that when he started out, he benefited directly from the free-flowing hacker culture of the time, but as Microsoft became ever more and more powerful, he systematically shut everyone else down, and stopped the party. So again we see licences as a defensive move against stupidity.

The Internet wouldn't be the way it is today without significant input from Microsoft. Many Open Source advocates would treat this as heresy, but anti-MS feeling in the Linux community detracts from real issues in other areas.



IP lockup stifles innovation What I think will eventually happen is that people will realise, *"Oh, the industry has more innovation when you keep a certain amount of fluidity!"* What happens is that when people lock up too much intellectual property – whether it's through patents or through copyright or through secret APIs or whatever method it is – you stop innovation, and the industry becomes stagnant. The best example of that is Microsoft. Microsoft has probably made more money from Open Source than any other organisation on the planet. But again, no one's telling that story. Why not? Because we define "Open Source" too narrowly.

But I look at the situation and say, *"Here's Microsoft, and it was totally stagnant. They were working in the labs, and what did they come up with? Microsoft Bob."* That was the next generation of computing. That was what we were all told: we've figured out what's going to come and be the next thing that will make you excited. And everybody replied: *"You've got to be kidding!"* And all we've got left now is that nasty little talking paperclip, right? What *really* saved their ass was Tim Berners-Lee putting some software out into the public domain back in 1992; they looked at that, and TCP/IP, which came out of DARPA and UC Berkeley, and all the free software that was happening there, and thought, *"Oh, cool, the Internet!"*

So Windows 95 was a direct result of Microsoft being able to capitalise on all the benefits of a huge swathe of technological stuff that had happened outside its hegemony in the Open Source community. In the end, what frustrates me so much about the Open Source community – as it stands at present – is that it's so focused on Linux. The Free Software movement in particular is so focused on replacing Microsoft as the challenge, yet they sound very surprised when I go, *"My God, the Internet! The Internet! It's something we all built together. It is the greatest triumph: it's ubiquitous, it's the killer app, it was built by an Open Source community!"*

Sure, there was some government funding, but we all helped put it together, and the killer app – the world wide web – was actually put in the public domain: you can't be more open than that.

Microsoft Bob: Probably not the future of graphical interfaces. As a forerunner to Windows 95, this result of millions of dollars of research only lasted a year... a whole 365 days too long, in many users' opinions.



More Open Source innovation And again, why isn't the Open Source community saying, "Hey, HTML – what a great success for us." How about JavaScript as well? They're just off the radar, because they don't have the political element: and that's stupid. We're not owning the fact that we created the revolution that dominates the industry: in marketing terms, we're basically still chasing Microsoft's tail lights; when in a technical sense Microsoft just implemented our stuff after chasing *our* tail lights – only we seem to be too shy to position it that way.

Look at *Apache*. Microsoft has spent ten years trying to beat *Apache* and it hasn't done it yet (judging by the market-share of Netcraft), and *Apache* just keeps climbing. I feel like that's a real success story there, like the Internet. The fact is that Microsoft built Windows 95 and Windows 98... all those generations of technology since Windows 95 were all an incorporation of stuff that came from outside Microsoft. That's a huge part of the innovation. Now with .NET, MS is saying, "Now we're going to try to do to that platform what we did with Windows – we're going to try to control it, close it down." Fortunately, the public – particularly the developer public – is hip to that. For example, when Microsoft came out with *Passport*, and said, "Wouldn't it be nice, we'll do it all," potential customers thought: "We've seen that movie, we didn't like the ending last time, so we're not going to watch the sequel!" I think that's a very interesting piece of the story.

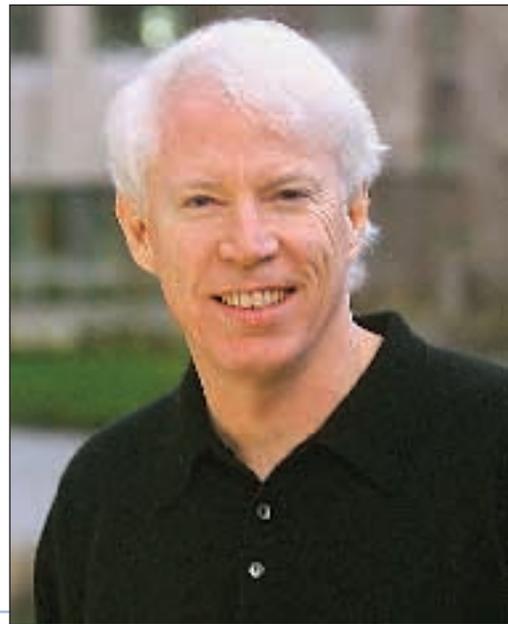
Preserving the culture I'm not saying Linux is unimportant – Linux is hugely important, as it's a big piece of the platform – but it's only a piece. A lot of people think they see, "Here's Windows, and Linux is playing catch-up". What I see is that there's a new platform: the Internet. Microsoft is trying to take that over, and we need to say, "We're already up here, guys!" How do we make sure that what we built stays ours – that

"LICENCES AREN'T EVEN A NECESSARY EVIL... THE OPEN SOURCE COMMUNITY IS MAKING TROUBLE FOR ITSELF BY TRYING TO CONTROL TOO MUCH"

they don't take it over? As opposed to thinking that the game is to get to parity with them, *they* are trying to get to parity with *us*, we're just not owning our successes.

Meanwhile, there's a new set of application providers building on top of what we have collectively built, and we're not spending enough time making them say, "OK, *this is our community – these are the rules we need to learn to play by!*" A lot of these next generation companies are going to say, "I want to be like Microsoft", whereas I want them to say, "Hey, I want to keep that culture that we came from." So it's real important to get people like Larry Page, Sergey Brin, and Jeff Bezos to say "Yeah, those are my forefathers – they taught me something, and I want to keep that in my DNA, because that's going to be good for me. Because if I don't do that what's going to happen to me is I'm going to close everything down and ten years from now someone's going

Jim Allchin, Group Vice President, Platforms Group, Microsoft – and visionary of a networked future.



to come along and do to me what we're doing to Microsoft." Anyway... I guess that's a very different perspective to what you'll usually get on Linux.

LXP: Earlier, you mentioned that you thought that licences weren't that important – they were just housekeeping. Does that mean you're not too worried about the recent announcements of problems between the *Apache* licence and the GPL, and also the X11 licence and the GPL?

TO: I haven't really spent a lot of time thinking about it, because I really think that licences aren't even a necessary evil; I feel that the whole Open Source community is making trouble for itself spending time trying to control too much. At the end of the day, Richard Stallman has just as much desire for control as Bill Gates does, and I don't identify with *either* of those extremes. I want to find that middle ground where we're focused on creating value, not managing to control what other people do – I don't *care* what other people do, as long as they don't try to control what *I* do. I want to see people stop writing licences that try to control what other people do. There may be specific reasons why people say, "Hey, I've really got to write this licence this way," but I think it's just a massive diversion of energy.

If I was Microsoft, I would just be so happy to see the way that people are focusing on licences – that's where MS keeps focusing everyone else's energy. But meanwhile, I meet with Jim Allchin (Group Vice President of Platforms Group at Microsoft) and he's *totally* onboard with this vision of a networked future, and that that's what MS wants to own. He literally said to me, "It's just like GUI – nobody owns it." And I thought, "What do you mean, nobody owns GUI?", but that was the vision they started from: there was this new platform – the Graphical User Interface – and *nobody* owned it. And now there's this new network thing, and nobody owns it – and they want to own it! And that's the fight for the future. Again, I feel like our fight should be to make sure that *nobody* owns that, and that we keep the culture where it remains in commons. That's way more important. 

COVER FEATURE **TIM O'REILLY**

« **A question of philosophy** I personally am a very strong *Apache*/BSD licence fan. People should be able to use whatever licence they want – if people feel strongly about it, they can put the software under the GPL, but I think it creates some huge problems even as it tries to solve others. To me what's important about Open Source is that we lower the barriers for people to innovate, to try things, and the GPL is so focused on being a defensive licence to make sure that people *don't* do something.

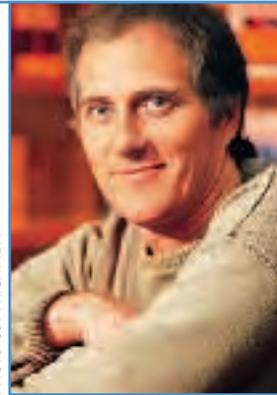


PHOTO COPYRIGHT IBM

“EDUCATE COMPANIES ABOUT THE VALUE OF HACKER CULTURE – IT’S ABOUT PUSHING THE ENVELOPE, IT’S THE FRONTIER WHERE NEW IDEAS COME FROM...”

TIM O'REILLY

I want to see licences that encourage people to *do* things. I can still remember my conversation with Bob Scheifler of X, from back when we were doing our X books and we were enhancing the MIT documentation, and that it was becoming free.

We got some flak about that, and I went and talked to Bob. I said, “*Hey, this is his software, his vision – how does he feel about it?*” That's what I've always tried to base my philosophy on: respect for the wishes of the creator of the material I want to re-use. Bob said, “*No, that's what I want. We're just trying to create stuff that people will take and run with, and they'll build on top of it.*” I love that philosophy: when you're giving a gift to the world and saying, “*If you*



Dave Stutz, Author of Microsoft's Shared Source .NET platform, Rotor; now full-time vintner.

want to commercialise it, great; if you want to keep it free, great – it's your choice.” Nothing is taken away if somebody does a proprietary fork – either it works and people say, “*Hey, that's good enough!*” and I don't care because they are good stewards of the software; or it just doesn't work.

And with *Apache*, what we've seen is that everybody who has tried to do a proprietary fork has failed; so, those licences I criticised earlier actually work reasonably well, and there's far less bullshit about “*You didn't quite do this, and you didn't quite do that.*” When I look at the GPL, I see it as an obstacle to my goals. I'd rather see people spending less time fussing about “*Is it compatible?*” or “*Can we re-use it?*”, and say “*Yeah, you can do whatever you want.*” To me, again, this has to do... well, there's this wonderful line from *Tao Teh Ching*, by the Chinese philosopher Lao Tze. He said: “*Losing the way of life, men rely on goodness; losing goodness, they rely on laws.*”

Why licenses? Licences are laws. Let's try to encourage people to do the right thing, as that's better. But even better than that is learning the way of life, which I think is just the science of it, understanding that openness works. And we can trust that openness works, because it really does. And what we have to do is teach people what works, and what I see working is people saying, “*Hey, I need to be open because it has these benefits, and over here I need to be closed because it has these benefits, and basically finding an appropriate balance.*”

On the other side, to the folks who fear that Open Source is an ‘intellectual property destroyer’, I like to point out that there are lots of opportunities for proprietary advantage – even in the most open systems. Let me give you a couple of examples. Take the open architecture of the IBM PC that I touched on earlier, which launched the commodity hardware revolution. I just point to the label on most PCs: ‘Intel Inside’. Here's this Open architecture of IBM PC, but somebody managed to create a proprietary spot in there – there are probably a bunch, but that's the best known. Now that we have the Open Internet, it could be branded ‘Cisco Inside’. When we move up to the software stack, there will be plenty of proprietary pieces hidden in plain sight. There's some right now, even in the Open Source Internet.

The business model for *BIND* was domain name registration, we just didn't realise it at the time and so it became separated and it's now a monopoly. That is basically worth hundreds of millions of dollars a year; it's a valuable property. There will be those kinds of things. If we think ahead we can anticipate and make sure they don't go to the wrong people. It would be nice if Paul Vixie (Author of several RFCs and well-known Unix system programs, among them *SENDS*, *proxynet*, *rtty* and *Vixie cron*; and founder of the Internet Software Consortium, which in 2004 was renamed Internet Systems Consortium) was making some of that money, and if it were a bit more competitive and a bit more open. But also, I think innovation has stopped in many of the consumer-facing applications; however, there's a lot of innovation going into building out this next generation infrastructure for computing. I think that what we're going to see is that we now have these new consumer-facing applications – the

Amazons, the eBays, the Googles – whatever it is that fills the application layer: and we're going to start weaving together an Internet operating system. And there are going to be a lot of pieces inside of that, owned by people like BEA (www.bea.com). People are going to figure out the hard parts of that problem, and eventually we're going to have utility grid computing, and there are going to be people who figured out the hard parts of that problem – I know of start-ups that are working on it right now.

Those guys are going to have deep intellectual property, as they are doing a lot of innovation and are going to end up controlling big pieces of the future. And one of them is eventually going to emerge and be the next Microsoft. Who knows? It could be anyone. But there are people who are building the insides of that global machine that we're now seeing the first prototype of in today's Internet.

LXP: Isn't the counter-argument that one of the biggest benefactors of Open Source software, through the BSD licence, is Microsoft?

TO: So?

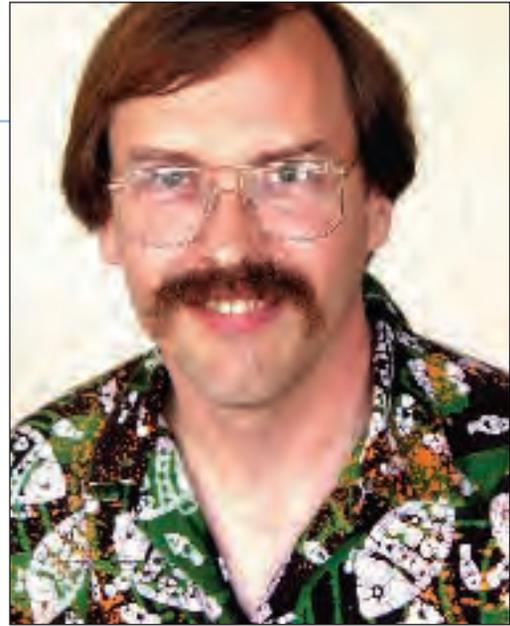
LXP: As you mentioned earlier, MS has created systems that are closed out of systems that were previously open and available to everyone.

TO: My reply to that is, "Yes, if you're Richard Stallman then absolutely that is a bad thing, because his goal is the creation of free software." He doesn't care whether that's better software – he has actually said, "That's not what I'm about." He's not about creating better software, he's not about creating more functionality for users – he's about creating Free software, which is a moral, ideological point of view. If Microsoft hadn't been able to take all that BSD-licensed code, the TCP/IP stuff, and we weren't able to have Microsoft helping to drive the commercial adoption of the Internet, we'd still be sitting here, like we were ten years back, with a bunch of geeks who have this wonderful Internet functionality but the average user probably would not have it.



Andy Hertzfeld ex-Apple full-timer, now at the helm of Chandler development.

Larry Wall: "Information doesn't want to be free; information wants to be valuable."



And yeah, maybe we'd still have our party, and maybe we would have marginally more success, but there would be hundreds of millions of people worldwide who wouldn't have those benefits. Take a look at some of the hacks that are being done. Here's one I picked up off Slashdot that I love: there's a wireless 'pony express' happening in Cambodia, where some guys have set up five motorcycles equipped with wireless, and they drive this route through rural Cambodia, and once a day email gets picked up. And there's a base station back at a school where they connect back up to the regular Internet – what a wonderful hack!

The hacker impulse – the freedom we really care about – will happen even on Windows. We just need to inject more freedom; we don't need total freedom. I look at that and think, "OK, so here are people in rural Cambodia who are able to the benefits (or, in some cases, they are probably getting spam – the drawbacks) of the Internet. And this thing that was created was partially done because, yeah, Microsoft got a lot of stuff out of the Open Source community, and they took it in, and what did we lose? We didn't lose any freedom. Yeah, some other people got proprietary software instead of nothing." The whole of the Internet was undeniably driven by commercial adoption by people like Microsoft. It would still be a geek thing if it weren't part of Windows. We'd still be a very distinct minority; there'd be 10 million people using the Internet, instead of hundreds of millions.

LXP: There are quite a few people in the Open Source community who couldn't see that as a disadvantage!

TO: No, but if you're looking at the creation of value – I see this as the natural evolution. The hackers are always going to move on. What I think really matters to the computer industry is that we need to educate companies – including Microsoft – about the value of hacker culture. Hacker culture is about pushing the envelope, about pushing the boundaries. Open Source is a facilitator for hacker culture, and openness of various types is a facilitator for hacker culture. And it's the frontier where new ideas come from, and we need to go for that frontier. ■■■

ORACLE

GET FLEXIBLE WITH APP SERVER 10G

In the second part of our focus on Oracle's 10g initiative, **PAUL HUDSON** investigates how you can consolidate your applications in one place and save money...

The web services revolution brought with it heterogeneity of software, distribution of resources, and a potential management headache. Yes, the numerous benefits of well-implemented web services more than make the case for its use, but it has added another thorn in the side of system administrators who really have better things to be doing with their time. Is micromanagement the future of computing? Is a unified data policy an unobtainable dream? Why can't the EE in J2EE stand for "Extremely Easy?"

Although Oracle is best known for its database system, it has more than one trick up its sleeve. *Application Server* has been going for some years, and continues to be immensely popular amongst those in the know. If you cower away from the term 'middleware', thinking it's a buzzword best left to your pointy-haired boss, let us dispel that myth: middleware is a key piece of the infrastructure in any successful company, and doesn't mean changing the way you work.

Being middleware, *Application Server 10g (AS10g)* is a little hard to define. Its uses are multifold, and this also makes it hard to define – it's not a web portal; it's not a wireless communication server, a *J2EE* server, or even a business intelligence system either. Instead, it's all four in one, with much more added. Crucially, this release adds full support for grid computing, building on top of the 'Zero Unplanned Downtime' featured in previous releases.

ALREADY USING BEA?

Oracle wants you to switch...

A year ago, Oracle launched its *BEA Switch & Save Program*, which is an aggressive competitor upgrade scheme designed to tempt existing BEA customers to switch to *Oracle Application Server*. And how tempting is it? Well, for every processor you have running BEA's *WebLogic*, Oracle will give you a free licence for *Application Server*.

Furthermore, if you show Oracle the invoice for your BEA support costs, Oracle will guarantee that you aren't charged more than half on all your switched processors.

That is, if BEA charged you \$50,000 a year,

Oracle will charge you no more than \$25,000 a year, guaranteed.

This all comes under Oracle's "*features BEA can't match, at a price they can't beat*" heading, designed to make plain the advantages – both in features and economy – in switching to *Application Server*. Although BEA has yet to fight back with a marketing plan of its own, it's only a matter of time. As always, the end-user is likely to be the biggest winner as the two heavyweights go head-to-head – look out for lower prices, better support, and more functionality soon!

BUSINESS INTELLIGENCE

Like 'middleware', 'business intelligence' is another term that has been bandied around to the point where it has little meaning. However, its original meaning was the ability to run arbitrary queries on data to enable ad hoc analysis in real time. *Application Server* tackles this problem in two ways. First, through *JDBC* it is able to connect to and query just about any data source you have available. As a result, your data can be analysed remotely irrespective of whether you're using *PostgreSQL*, *Oracle*, or even *Microsoft SQL Server*. Second, *AS10g* can generate reports in a variety of formats that you can subsequently use either as part of your data analysis, or, for maximum flexibility, use as part of your personalised web marketing for customers.

A different kind of business intelligence is how *AS10g* automatically responds to each user's access requirements depending on their device or their location. So far we've only looked at how *AS10g* works with the web, but it works equally well in other media – it delivers content to devices, irrespective of whether they are laptops, PDAs, or even Java-enabled mobile phones. If your phone isn't Java enabled, applications can even take advantage of voice recognition to provide some degree of control. Pushed content can include SMS messages to mobile phones, emails, or automated voice calls; and *Application Server* can be configured to vary its settings depending on the access location and situation of the user – for example, it might send little or no media when a GPRS connection is being used.

WEB SERVICES

Whether or not you already have web services in use, *AS10g* can nearly always lend a helping hand in the administration. If you already have services in place, the chances are they will all work out of the box with *AS10g* because it fully supports

the WS-I standard. Furthermore, thanks to *AS10g*'s abilities to communicate with such a range of other packages, you'll likely find your web service can be extended a lot further – even up to communicating with older, mainframe applications if necessary. This is boosted by the fact that *Application Server* includes a web cache for buffering popular applications. This has a major effect on performance, and, perhaps more importantly, the caching functionality is available as standard that reduces the need for administration.

J2EE

For some time now, Oracle has touted its *J2EE* server as the "lightest, fastest, and most reliable on the market", which, in a market with such serious contenders as *BEA WebLogic* and *IBM WebSphere*, is quite a boast. However, numerous benchmarks have shown *AS10g* leaving its competitors in the dust, which adds a degree of credence to the claims. As it is supported by every major Java IDE, it's easy to develop specifically for the system, or even to port from *WebLogic* or *WebSphere*. Furthermore, as Java is so portable, there is little opportunity for any vendor to lock-in customers to a specific solution.

With 10g ushering in grid computing, it will be interesting to see how Oracle exploits Java's inherent cross-platform nature when applications are distributed to varying hardware and operating systems. One sure bet is that Oracle will be amongst the first to implement *J2EE 1.4* now that it is ratified.

RELIABILITY AND SCALABILITY

If there's one thing that grid computing means to the masses, it's inherently better reliability and scalability. As we saw last month, you can add computers to a grid simply by plugging them in, and it will easily be more reliable than having the same number of computers working in discrete clusters.

As Oracle already had a legendary level of reliability, it's unsurprising that all the previous functionality for working with Real Application Clusters is still available if you want to use it. Although all the grid computing functionality is there and ready to use, old favourites such as Transparent Application Failover and Fast Start Fault Recovery Architecture remain intact. Furthermore, you can roll out new applications through *AS10g* without needing to restart the server, further ensuring maximum uptime.

Switching to the grid paradigm makes scalability almost a non-issue. Similar to the 10g database from Oracle, *AS10g* automatically administers computers within a grid, which gives 'plug-and-play' a whole new meaning. Also similar is that commodity hardware – even if it was considered top-of-the-range two years ago – has a new lease of life ahead of it as part of a computing grid.

ANY ROOM FOR STANDARDS?

Over the past few years, Oracle has gone through a massive transformation that has seen it move from being seen as a lumbering behemoth, to a leaner, more flexible company actively working to extend its products according to customer needs. Thankfully, a large part of this process has seen Oracle adopting various computing standards, not the least of which is Linux itself, and making them the heart of their strategies.

"YOU CAN ROLL OUT NEW APPS THROUGH AS10G WITHOUT THE NEED TO RESTART THE SERVER, FURTHER ENSURING MAXIMUM UPTIME..."

WHO'S USING APPLICATION SERVER?

It reads like a Who's Who of successful businesses...

Bank of Ireland, Barclays Bank, Boeing, CERN, DC Tech, Digital River, Honeywell, Lufthansa, Unocal, the US Army, the US Geological Survey, Verizon, and many thousands of other companies run *Application Server*, and for a variety of reasons: Boeing uses it to analyse and optimise its manufacturing process, Verizon uses it to provide salespeople wireless access to its corporate sales portal, and Lufthansa uses the business intelligence capabilities to analyse feedback and improve customer service.

The success of Oracle's *J2EE* efforts is self-evident, and the sheer fact that the majority of the web services discussion revolves around either Oracle, BEA, or IBM says a lot for the success of enterprise Java as a whole. In all the *AS10g* diagrams, *J2EE* sits prominently at the core, usually much larger than even database access – Oracle clearly sees *J2EE* as a key driver in both reliable web servicing and now grid computing also.

With over 16,000 companies already using *Application Server* – as well as it being the fastest-growing product in the market – *AS10g* is undoubtedly a big step forward for the product. 16 of the US's Fortune 20 use *Application Server* already, which makes this a pivotal market for Oracle, and one its competitors would be very happy to snap up if possible. However, it's quite clear now that grid computing isn't just vapourware – or worse, a buzzword – it's already being used to tackle high-end computing problems, and Oracle has really got the jump on other vendors in the market.

As you can see, it's pretty hard to nail down *AS10g* as just one application. Behind the scenes, there are at least twelve parts working together as one that previously would have been separate products. The upside to this is that as well as solving your immediate problem – which might be scalability, the need for Java, a rollout to mobile devices, etc – you'll also find it gives you a lot of functionality previously unavailable, and perhaps even new unthought-of ways of doing things. The drive towards flexible business isn't going to stop any time soon, or potentially ever, and with software like this, Oracle is making it clear that it aims to promote software consolidation as much as hardware consolidation – it's no surprise to hear IDC predicting the application server market to top US\$4billion by 2006.

In the final part of our look at Oracle's 10g launch, next month we'll be seeing how *Enterprise Manager* shapes up in its new 10g clothing, and also how *Collaboration Suite* may well be the dark horse in Oracle's product-line. ■■■

In a highly competitive market, you have to not only promote your own strengths, but also capitalise on competitors' weaknesses.

