

LINUX

FORMAT

THE UK'S BEST-SELLING LINUX MAGAZINE!



MANDRAKE 9.2

★ Better than ever ★ Install from our discs ★ Upgrade earlier versions ★ Latest software included

USER GUIDE STARTS p52

SHARQSTATION

Xinit's first contender for the Linux workstation market on test **p24**



LINUX EXPO 2003!

The news, products and gossip from the UK's No 1 event **p13**

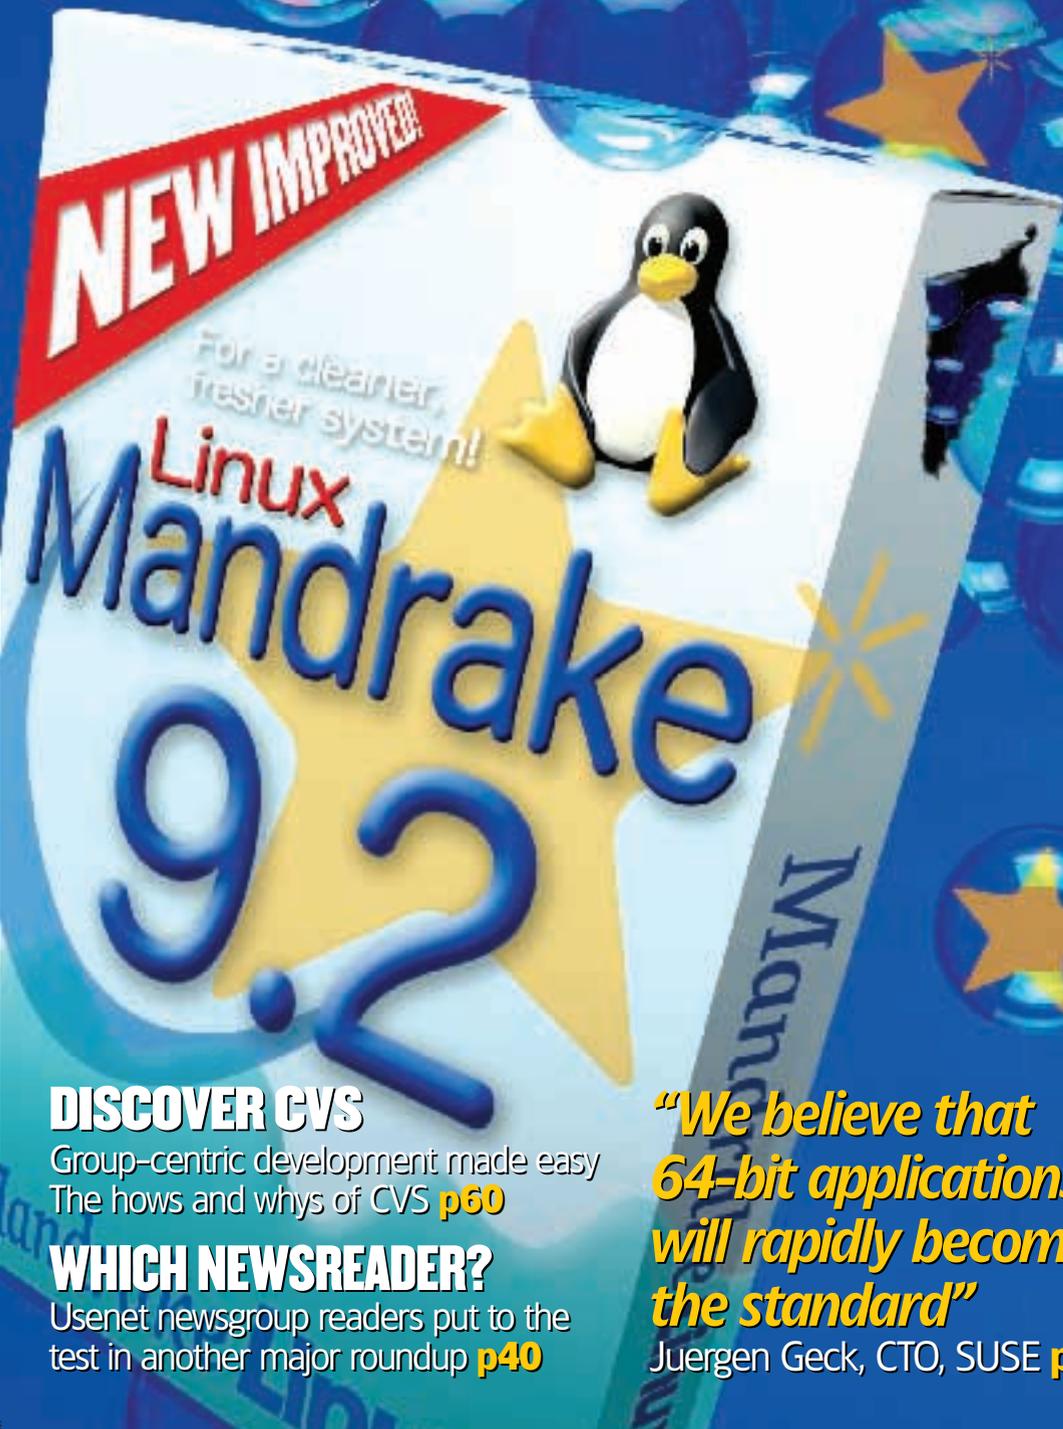
SUSE 9.0

A new look, a new logo and a fantastic new distro **p32**

LAZARUS – FREE PASCAL MADE EVEN EASIER? p30

FIND OUT HOW TO...

- Set up basic networking **p66**
- Write a compiler **p70**
- Make games in Blender **p74**
- Customise The GIMP **p78**
- Use sockets with PHP **p82**
- Build a firewall **p86**



DISCOVER CVS

Group-centric development made easy
The hows and whys of CVS **p60**

WHICH NEWSREADER?

Usenet newsgroup readers put to the test in another major roundup **p40**

"We believe that 64-bit applications will rapidly become the standard"

Juergen Geck, CTO, SUSE **p8**

CD issue also available Printed in the UK
LXF47 DECEMBER 2003 £6.49

 Media with passion 9 771473 968012 1 2
 (£6.49 outside UK and Republic of Ireland)

www.linuxformat.co.uk

Ready for anything?

While waiting to get a word in edgeways at the 'Great Linux Debate' at this year's Linux Expo UK (see page 13 for a full report), I found myself with a few spare clock cycles to ponder the questions that the audience were asking and the replies given to them by the luminaries on the panel. There were questions about the desktop, questions about Enterprise, questions about small businesses. There may well have been a question about embedded devices that I missed while pouring myself another shot of, er, water. The thing that seemed to get lost in the detail a bit was that Linux seems to be eminently capable of tackling any area of computing, from the very tiny to the very large. All it takes is someone to want to solve a particular problem and a bit of effort.

By some means or another, Microsoft managed to bag a place on the panel in the form of Bradley Tipping (brave man, and a nice enough chap, actually), who was quick to point out that businesses like IBM and HP weren't really doing

things for free, out of some sense of altruism, but only invested in areas where their business needs and those of the Linux community happily coincided. A very good point I thought. However, (and I may be wrong here, so please write in and correct me if so), I think that the interests of the IT industry and users coincide a great deal when we are talking about standards-based technology.

There may well be a case for proprietary software in business, there may even be a good case for software patents (though I haven't yet seen one), but there is no point whatsoever in closed standards. Who loses out if enabling protocols or file formats are – by means of intellectual property rights – the preserve of a single organisation? Everybody. The reason why Linux has been so successful in such a short time owes more to the free and open nature of its development than to the code itself. Once the world of IT catches up with the rest of science and realises this is an idea that isn't going to go away, perhaps we can get back to making progress.

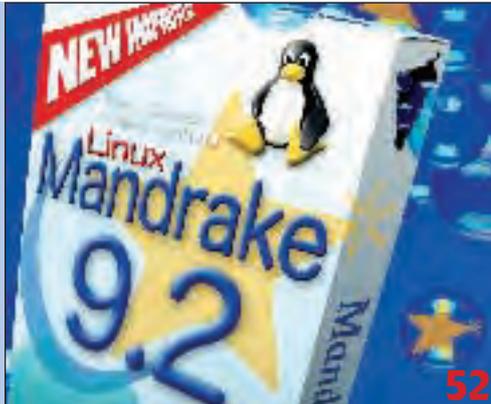


Nick Veitch EDITOR

New, improved and on your coversdiscs! Our extensive guide starts here... p52

What happens when 6,000+ Linux fans gather in the same place? Linux Expo UK 2003 is what! p13

A new look and a new logo aren't the only changes in SUSE 9.0 p32



MEET LINUX FORMAT'S TEAM OF WRITERS...



Andrew Channelle
Our newshound and Linux beginners' best friend, Andy also wrote the show report for Linux Expo UK 2003.

Richard Drummond
His greatest regret this month is not making it to the UK Expo for some proper beer!



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

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



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

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



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

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



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

Warren Brown
A Linux-loving IT pro all the way from South Africa? It's "as true as Bob," as they say in those parts, apparently!

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.

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

LXF47 December 2003

Contents



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



COVER HIGHLIGHTS

Mandrake Linux 9.2

52

The latest iteration is better than ever! Upgrade or install from coverdiscs with eight pages of help from our experts!

13 Linux Expo UK 2003 – Olympia 2

From the .ORGs to the corporates, fun was had by all!

24 Xinit SharqStation

Big hitter from server scene enters workstation market

30 Lazarus

A contender to Kylix, or just a free imitation?

32 SUSE 9.0

Supports more Winmodems than any other distro!



“One quickly discovers that these folks are ‘normal people’ with other interests and opinions unrelated to Linux...”

09 Hoyt Duff spends ‘a week with geeks!’

“When the site is ready to go live, I can easily use CVS to move my entire website to a new server with a single command.”

60 Jono Bacon sings the praises of CVS

LINUXPRO

Covalent worry-free web serving **David Valentine** IBM's Linux strategy **Thin Clients** Wyse's vision of the future **Sharqstore** Affordable NAS storage **Nessus** Scan and harden your servers **Hosting** Spotlight on Rackspace



27 MTS: theoretically a great idea, but what's it really like in practice?

06 News

Analysis of the latest Linux stories

18 Mailserver

Share your opinion with our readers

24 Xinit Sharpstation review

Serial ATA storage on the desktop!

27 MT Studio review

Rapid application development

32 SUSE 9.0 review

Latest distro put through its paces

34 Linux-compatible peripherals

Steering wheels and joysticks

36 Book reviews

The best treeware that's out there

40 Roundup: Newsreaders

Newsgroup and IRC readers rated

47 HotPicks

All the latest Open Source software

60 What On Earth... is CVS?

Committing changes to source code

66 Beginners' tutorial

Networking with Linux

70 Compiler writing tutorial

Loops and increments revisited

74 Blender tutorial

Introducing Blender's Game Engine

78 The GIMP tutorial

Customising patterns and brushes

82 Practical PHP tutorial

Revolutionise your code with sockets

86 Server School

Firewalls and what they do for you

90 Scanners tutorial

Tame that troublesome flatbed

92 Answers

Your queries answered

102 User Groups

Local support and fun

108 OFFER: 15% off MandrakeClub

Sign up and save money with LXF!

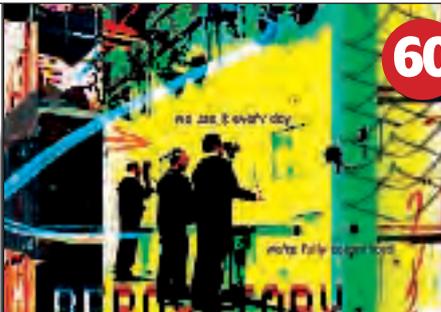
24



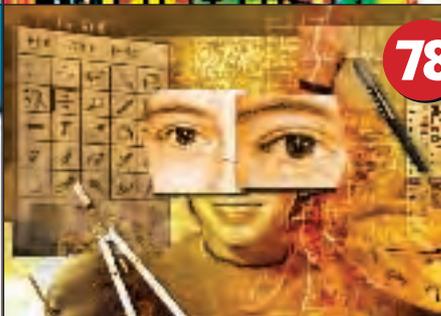
40



60



78



COVERDISCS

A DVD or 3 CDs packed full of Linux goodies **52** and **106**



CDS 1, 2 AND 3

Mandrake 9.2 the latest version of this popular distro completely fills the first two discs, with extra programs on disc three as well; **OpenOffice.org** latest iteration of Linux's most popular productivity suite; **GNOME 2.4** complete newest release of this capable desktop environment



DVD

OpenOffice.org PowerPC version; **MailScanner** kill spam dead; **XFCE** lightweight alternative to KDE or GNOME; **LDP** Linux Documentation Project; **PuppyLinux** very compact distro that fits on a USB keyring!

fphoto share image albums on the Web

Please read the coverdisc instructions on pages 52 and 109 before installing from coverdiscs!



SAVE MONEY!

SUBSCRIBE TO LINUX FORMAT

See page 100 or phone 0870 4448645



Newsdesk

- SUSE 9 launched ● Software patents ● Telewest wireless broadband ● SGI Universal storage?
- Opera on set-top boxes ● .NET Mono ● Linux meteorology ● SCO News ● Windrivers migration

PRODUCTIVITY SUITE

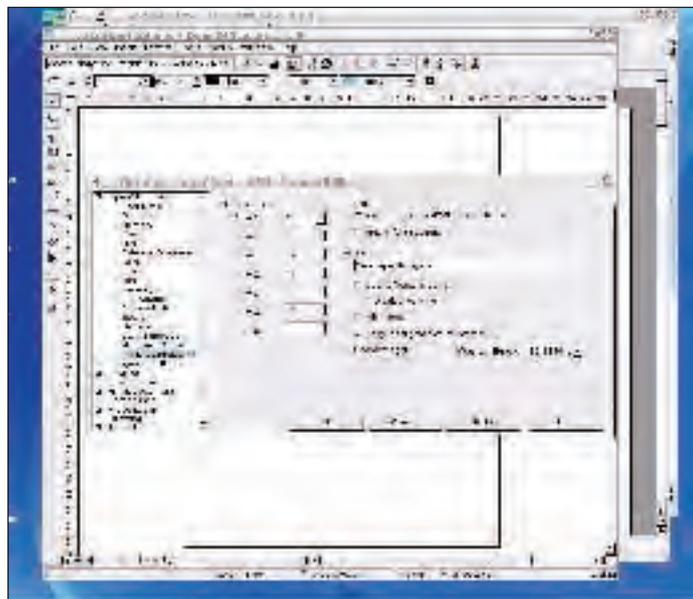
OpenOffice.org overhaul

The Autumnal release schedule is once again becoming a little overcrowded, with new versions and distributions from SUSE, Mandrake, Red Hat, *Samba* and *OpenOffice.org*. *Samba* and *OpenOffice.org* are being hyped, and not by their own developers, as the applications that could once and for all give Linux the foothold it needs to achieve significant growth in Enterprise.

OpenOffice.org 1.1 includes a wide range of refinements and feature additions, chief among them better support for the latest Microsoft file formats and the ability to output both Adobe's Portable Document Format (PDF) and Macromedia's Flash (SWF)

files. PDF support, in particular, is regarded as a key feature as the format gains ground as a *de-facto* (if not official) standard for distributing printed documents on the Internet.

Sun has already released its proprietary suite which polishes the *OpenOffice.org* interface and adds a few proprietary bells and whistles. Curtis Sasaki, Sun's VP of Desktop Solutions, said the release was once again a great tribute to the core *OOo* team and the community of developers and testers that has built up around the project. "The release of Version 1.1 of *OpenOffice.org* is a major achievement for the *OpenOffice.org* community. With the successful inaugural *OpenOffice.org* conference



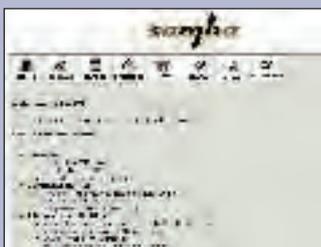
OpenOffice.org has had a comprehensive overhaul – see the LXF coverdiscs.

SAMBA 3.0 OPENS WINDOWS FURTHER

Cross-OS compatibility

While *OpenOffice.org* makes waves on the desktop, the latest revision of *Samba*, the Windows file/printer server solution, has gained a number of features designed to make building and maintaining multi-platform networks even easier. One of the most significant additions to *Samba 3.0* is the first Open Source implementation of Windows NT Primary and Backup Domain Controller. This allows users to 'transparently migrate their existing Windows NT domains to Samba while keeping their existing user and group account databases'. *Samba* is also now compliant with Microsoft's Active Directory, Kerberos 5 authentication, SMB signing and SCHANNEL security for secure remote procedure calls.

What this all means is that a Linux/*Samba* solution can be essentially in as a



Samba now boasts better Windows integration.

replacement for an NT file/print server with little difficulty and no client licenses.

The package also ships with an electronic copy of O'Reilly's *Using Samba* book (ISBN 1-5659-2449-5), which has been released under the GNU Free Document License.

www.samba.org

earlier in the year, the energy in the community drove many corporate and individual contributors to work together with Sun to create the most stable and powerful release of the *OpenOffice.org* application so far.

The suite is initially available for Linux, Window and Solaris (SPARC) with versions for BSD, OS X and x86 Solaris due before the end of the year.

Sasaki added: "New features like the native PDF export and the support for assistive technologies will accelerate the fast growing, worldwide adoption of *OpenOffice.org* and *StarOffice* as well as more secure, alternative desktop environments."

In fact a recent study by The Jupiter Group suggested that small to medium enterprises (SMEs), ie those with less

than 1000 employees, are taking to *OpenOffice.org* in quite significant numbers. The upstart productivity suite is said to have achieved a market share of six per cent, but report author Joe Wilcox said there is still work to be done, especially in the area of file format support. "They need to develop better translators so they can actually read or produce *Microsoft Office's* proprietary formats, because that's what's most widely used out there."

The report also claims that, in the USA, as much as 19 percent of small businesses are using Linux on the desktop and 26 per cent on the server, with robust, inexpensive applications such as *OpenOffice.org* on Windows paving the way for wider adoption of Linux in the future.

AERODYNAMICS & TUNING

Linux in pole position



Real-life Tux Racer – it's Linux that keeping this car on the ground!

After sponsoring the BMW

WilliamsF1 Team for four years, Hewlett Packard has furnished the team with a state of the art Linux cluster designed to "boost aerodynamic innovation and deliver the most intense analysis ever deployed in motor racing"

BMW Williams' new supercomputer is built on HP ProLiant servers and will be instrumental in the fine-tuning of the team's current crop of cars and, more importantly, will help model the aerodynamics of the new car design which will be unveiled in time for the opening of new F1 season in Melbourne, Australia on March 7 2004.

HP and BMW Williams have been working together over the 2003 season to improve simulation capabilities, HP claims, by a factor of three. In the real world this means that the team has been able to shave

days off of the design and test process, especially in the field of Computational Fluid Dynamics. Nathan Eagle, senior CFD aerodynamicist at WilliamsF1 says the deployment of the teams new Linux cluster doesn't just help streamline the testing process, it has also 'significantly contributed to recent performance increases.' "It allows us to cut design, development and testing time by half, while, at the same time, to provide us with more capacity to experiment new car design concepts," he says. HP's director of corporate sponsorship, Andrew Collis, says that the technology helping to keep F1 cars on the ground in 2004 will inevitably find its way into the consumer market and aerospace industries. "The benefits this solution offers are by no means confined to the Motor Sports field."

HELPEXPLORER

Help for Windows help files

KAMA software has launched a new application which allows Linux developers to port help files that have been created in the WinHelp format for Windows applications, across to Linux with very little effort.

In fact, *HelpExplorer* allows developers to use the same help files with no editing or reprogramming. In addition to a pair of search features, covering index and key phrase search, *HelpExplorer* has an extensible, multi-

window design system complete with multiple history support.

A single user license costs US\$69, while a site license – allowing you to deploy the system on every computer in an organisation – costs US\$499. A full deployment license, for developers intending on including the help system in distributable applications, comes in at US\$749. A trial version of the software is available at www.kamasoftware.com.

NEWSBYTES

■ **Red Hat's** finances are looking a little better after posting decent results for the second quarter of 2003. The company increased revenue to US\$28.8 million, a sequential increase of six per cent over the first quarter, and achieved a net income of US\$3.3 million, compared to a net loss of US\$1.9 million in the same quarter last year. Red Hat CFO Kevin Thompson stated that the results reflected a strong customer demand for Red Hat Enterprise Linux Solutions. "The consistent improvement in our gross margins over the last three quarters to 72% for the second quarter shows the significant scalability of our subscription business model," he said.

■ Though still only a fledgling project, *Scribus* has competition in the Open Source DTP stakes. **Passepartout** (pictured below) is a native GTK application. It is very early days for the application yet, version 0.2 has just been packaged up for Gentoo, but already it is showing signs of promise.



■ **Mandrake** recently inked a deal with small form motherboard and PC manufacturer Shuttle to pre-install Mandrake Linux on its range of XPC network appliances. The company's 'barebones' network appliance range is aimed at small businesses needing entry-level servers or workstations. "Shuttle's Network Appliance XPC product line is the perfect platform for delivering and using Linux," said Mandrake's Regis Wira. "Putting these two best-of-class products together is a natural evolution of the desktop computing. This synergy offers a stunning new level of performance and value."

■ **Penguin Computing** is shipping two- and four-processor Linux machines utilising AMD's Opteron. The Altus 4200 uses Opteron 800 Series processors and supports up to 32GB RAM and 584GB of storage on four hot-swappable SCSI drives and is aimed at customers wanting high performance, symmetric multiprocessing with support for both 32- and 64-bit applications.

■ Audi has migrated its metal casting simulation and visualisation systems to a Linux cluster from Linux Network using 34 Intel Xeon processors. Updating from its old RISC-based system has not only saved Audi money, but has improved performance by a factor of ten. Simulation that previously took two weeks can now be completed in two days.

Jono Bacon

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



COMMENT

Linux & tech culture?

“ A few years back, the common computer user who generally hated computers and having to use them, typically used an OS such as Windows or maybe (particularly arty types) a Mac. These OSes capitalised on the easy-to-use factor and hid most of the technology away behind icons and frustrating paper clips. In recent times however, I suspect Linux's explosion onto the scene is fundamentally changing user attitudes to technology.

Ask yourself the question "why do people like computers?" With a few exceptions, the answer is that most people like to tinker with things, and computers are like an everlasting gobstopper: they have endless possibilities. I suspect that when people learned there was an OS that did not have the barrier to the technology, it attracted many users. I know people who are not beard-growing techies, but just like playing with their computers. Linux offers the technology for the hardcore hackers as well as the point-and-click generation with the variation in distributions.

Operating Systems need to be flexible. It is the OS that provides the limitations on a computer in terms of software, and Linux has provided a means to show people what really goes on behind the scenes. Five years ago, the typical technically interested computer user had a few networked Microsoft machines running Windows 98. Now the same kind of person has a Linux network with Smoothwall firewall, mail/web/cvs and other servers and has everything controlled by a web interface. The culture of computing is certainly changing, and Linux has opened bonnet. ”

NEW FEATURES & FUNCTIONALITY

Update your distribution

SUSE is adding User Mode Linux, Mandrake is adding adverts for download only users, and Red Hat is completely rethinking its desktop business in a bid to remain in the black. Yes, backup tools at the ready: it is upgrade time!

As well as changing the case of the U in its name, SUSE has announced the availability of SUSE 9 in the usual range of versions. Built on a custom 2.4 kernel, SUSE's developers claim to have backported many of the big advances in the latest kernel iteration. For more adventurous users, there is a testing 2.6 installation routine – which is a little more selective about the hardware it can use – and, in a special release, native support for AMD's 64-bit Athlon.

The company's chief technology officer Juergen Geck said he expected 64-bit computing to take off in a big way, but the Athlon's ability to run 32-bit applications provided the perfect opportunity for buyers to migrate their

systems in their own time. "We believe that 64-bit applications will rapidly become the standard, and that's why SUSE had an operating system for the AMD Athlon 64 processor available at the time of its introduction," he said. "With SUSE Linux 9.0 for AMD64, customers can run their 32-bit programs until they are ready to take full advantage of 64-bit architectures."

Other advances in the latest revision include full support for resizing NTFS partitions, access to a wider range of so-called WinModems that are often included in inexpensive desktop and laptop computers, a new YaST module for easier Samba configuration and an integrated build of User Mode Linux, which allows users to run many different instances of Linux on the same machine concurrently.

Mandrake 9.2 has been under heavy beta testing over the past few months, and the company has come up with what it thinks is a solution to the perennial problem of allowing access to



SUSE 9 is on sale from October 24th 2003.

the distribution to those unable or unwilling to support the company financially: advertising.

The sales pitch highlighted the millions of downloads and sales every year that "offered an opportunity to dramatically increase" visibility for companies and products with a

demonstrable commitment to the Linux/Open Source community. There are four spaces up for rent in the download edition. Installation, default web page, web browser bookmarks and screensavers.

Finally, Red Hat has divorced management of its consumer focused

Linux Web Watch/



FFII has a wealth of useful links.



EFF – an opinion for every occasion.



Owned slaves, but against patents!



Know your adversary... the BSA.

Software Patents

Software patents, Free Software advocates claim, should be concern not just for application and OS developers, but also users.

Recently the European Parliament passed a directive enabling the enforcement of patents across the community. There are a number of organisations attempting to educate legislators on the pros and cons of adopting the system.

<http://swpat.ffii.org/>, which is part of the EuroLinux protest effort, has a large selection of links to resources

including an extensive recommended reading section. The only problem with this site, is that it could do with a little updating to take into account the latest developments.

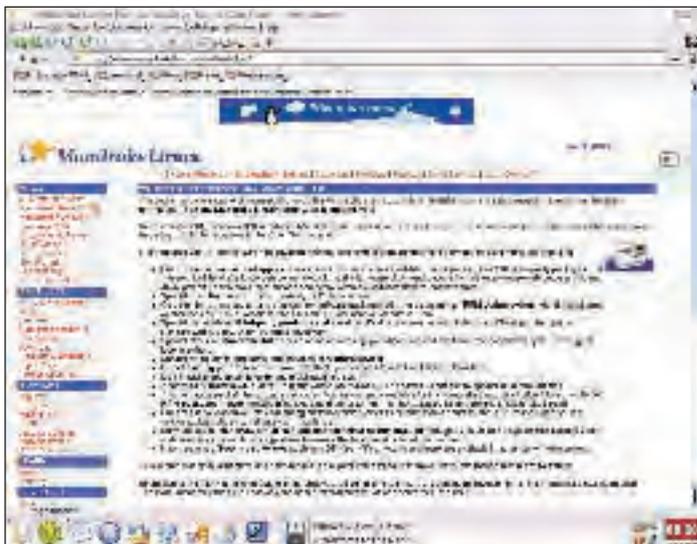
The Electronic Frontier Foundation is active in many areas of online life – including file sharing, fair use in copyright, digital rights management and software patents. This is a

comprehensive resource for anyone interested in the policing of the Internet.

For a little historical context on the situation in the USA, which already has software patent laws on the books, see <http://odur.let.rug.nl/~usa/P/tj3/writings/brf/jeff220.htm>, which is an online archive of Thomas "no patents on ideas" Jefferson. What you'll take from this site is the extent to which

copyright laws have been changed in favour of producers.

A contrary view can be found at the site of the Business Software Alliance (www.bsa.org). This organisation was set up by software vendors to represent them in government circles and work on enforcing current laws with regard to piracy and licensing compliance.



Mandrake club users get early access to the latest version, non-paying downloaders can expect the introduction of Ads in Mandrake 9.2. For reduced-price membership, exclusive to LXF readers, see page 108!

distribution and merged it with the Fedora project. Red Hat Enterprise Linux, which is currently bringing in the money for the company, is unaffected by the change. While Red Hat will maintain "editorial control" over Fedora, there will be some big changes in the way the desktop product evolves and its content, which is said to be aiming for more bleeding edge territory.

The first major change is that Fedora/Red Hat will push out updates three or four times a year and full support will lapse after just three months. The latest beta, codenamed Severn, includes a number of advances including an Apt-capable version of RPM which should make it easier for users who 'update early/update often' to stay with the game.

BROADBAND

Cheap wireless pack



Telewest's wireless kit is based on standard Netgear components.

Telewest has launched a new product under its Blueyonder name which uses the modem built into the cableco's set top box (STB) and some Netgear kit to provide 'one to one' wireless broadband access. The service, which costs £35, plus £25

per month for the 512KBps version, provides users with a Netgear ME102 access point to plug into the STB and either a USB or PCMCIA wireless NIC, both of which can be configured and used with most recent Linux distros. www.telewest.co.uk

NEWSBYTES

■ China's **RedFlag Linux** is reportedly preparing to release an English language version of its distribution for the first time. The native version of RedFlag Linux Desktop 4.0 was released in July, and the imminent English version is intended to give the company a global profile. RedFlag recently signed a partnership deal with Hewlett Packard to begin promoting the distribution initially in China and later around the world.

■ **IBM** has also signed up with RedFlag to bundle DB2 Express with the distribution.

■ Lindows.com has signed a deal to pre-install LindowsOS on a range of hard drives from **Seagate**. The intention is to allow OEM system builders to cut costs, while still offering a fully-featured computer.

■ **Lindows** boss Michael Roberts has once again incurred the wrath of Microsoft with a new website called MSFreePC.com which was set up to help Californians claim their portion of the payout from a recent class action suit that Microsoft lost. Roberts claims he is just providing a service to the public and making claiming a more simple process.

■ Austria's **antitachyon** - Manalo & Willner OEG, has released the fourth version of its *Server optimised Linux (SoL)* product. The latest release has a new installer and installation of the *Zope* application server and an enhanced XML-based boot process to improve loading times. <http://www.sol-linux.com>

■ The Linux trend continues across Europe with news that the city of **Vienna** is mulling over migrating 15,000 local government desktops to Linux. The city will make the decision in early 2004 with implementation expected by 2007. Sun's Mad Hatter project has been rechristened the Sun Java Desktop System and is built along the lines of Ximian's Desktop product.

■ At a joint conference with IBM, Microsoft chairman **Bill Gates** committed his company to maintaining a royalty-free system for its proprietary web services technologies. Gates enjoyed (maybe) a demonstration of web services using Linux and Netscape software to demonstrate this new inclusive strategy.

■ The latest **Xbox update**, said by some to be installed whether requested or not, fixes the security holes in the console's Dashboard which allows Linux to be installed on an unmodified device.

■ A Linux cluster at the Australian Centre for Advanced Computing and Communications has broken the teraflop barrier - making it the **world's fastest supercomputer**.

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

A week with geeks

“ Many of us labour away in solitary, having little actual contact with other geeks, save a few close mates and perhaps a few LUG members. Fanatical Linux users - and computing types generally - are not perceived by the public to have an excess of social skills or even a need for frequent human contact.

How wrong they are! I spent a week with geeks - not only the regular ones, but a few uber-geeks as well - in a fun and relaxing social setting: the 2003 Linux Lunacy cruise. Sailing aboard the Holland American Lines' magnificent *Amsterdam* cruise ship in Alaskan waters, 120 geeks attend seminars moderated by uber-geeks such as Linus Torvalds, Bruce Perens, David Axmark, Keith Packard, Ted Ts'o and others. Not only was there interaction between these luminaries and we mere geeks during the seminar sessions, but at breakfast, dinner and the occasional pub crawl.

What one quickly discovers is that these folks are "normal people" as well, with other interests and opinions unrelated to Linux and they are plagued by the same problems as we. Not only are they mortal, but more interesting than you might have imagined.

I also had some fascinating interaction with the geek wives. These women are intelligent, opinionated, graciously social, and have diverse interests. While most are not Linux users (many use MS Windows), they take pride in their husband's accomplishments. Aware that their husbands are different than most other men, they are among their most zealous supporters. We should all be as fortunate.

SHARED DATA SOURCES

Universal storage enabled by SGI?

Silicon Graphics Inc (SGI) is touting its third-generation *InfiniteStorage* system as the solution to the age-old problem of accessing shared resources in a multi-platform world. *InfiniteStorage* links in with SAN switching and the company's CXFS clustered file system to allow Windows, Linux and Unix servers to access shared data sources. Out of the box, the solution supports most Unix variants including AIX and Solaris, 32- and 64-bit Linux and Windows NT/2000. Support for Windows XP Professional and Apple's OS X is in the pipeline.

SGI spokesman Ajay Anand said the new product line would bring SAN-NAS storage capabilities, usually reserved for supercomputer environments, into the sphere of enterprise use. Storage options in the product line scale from 1TB up to a massive 150TB.

One of the key markets for this technology, SGI says, is the visual effects industry, where enormous amounts of data need to be managed. The Orphanage, the effects house

that handled the CGI work on the *Spy Kids* and *Charlie Angels: Full Throttle* movies, was one of the first customers for the latest upgrade. IT Director Nicholas McDowell said that they had only been living with the system for a few months, but that its potential was 'fantastic'. "It gives us the ability to scale, in both data and in serving the data. As we add users and render boxes, we can just add additional file servers that are accessing exactly the same data"



Orphanage clients include Nike and Taco Bell...

SGI intends to make multi-platform networking less of a chore.



Embedded Linux News

- **Operasoft** has begun a major marketing push to convince set top box manufacturers to build their interfaces around Linux and the *Opera* web browser. The Norwegian company says such a solution would help dramatically cut costs for STB builders while adding cutting edge facilities. Opera has previously sold browser technology for set top boxes, but now claims that developments in the browser will allow it to interact with the electronic programme guides (EPG) that navigate and present menus and other information to the TV viewer. Lars Boilesen, Opera's executive vice president for sales and distribution, said that the Linux and *Opera* solution would help system-builders hit the sub-\$100 per device target.

- **NEC's** efforts to develop high volume IP packet transmission capabilities for video streaming and Global Positioning System (GPS) services on 3G mobile phone

handsets has borne fruit with the release of a new telecoms package built on Carrier Grade Linux.

- The **Open Source Development Labs** project to develop a robust consumer entertainment Linux OS has been boosted by the recent addition of Japan's telecommunications giant **NTT**. "We think that there are several challenges for open-source software in mission-critical systems," Kou Miyake, CEO of NTT Data Intellilink, said. "We believe that OSDL provides a unique bridge between users, vendors and open-source communities."

- Linux as a mobile phone OS has received a boost with the announcement from **Openwave** systems - which designs mobile web and email access software - will be supporting the platform. The release of Openwave's Phone Suite software is thought to be a reaction to Motorola's strong commitment to Linux.

MORE EMBEDDED LINUX NEWS

Windrivers' change of heart

The world's largest embedded software developer has had a dramatic change of heart with regards to Linux. Windriver had previously snubbed Linux development, but after pressure from customers the company has announced a step-by-step approach to Linux migration, starting with its \$4,000 *visionProbe II* hardware tool.

Senior Vice President Dave Fraser said the company had taken a pretty hard line with the Open Source operating system but that "after much analysis, the conclusion is that there's a business opportunity, that Linux is a massive force, and it's here to stay in the embedded market."

Latest figures suggest that Windriver



The world's largest embedded developer embraces Linux... at last!

dominates the embedded market with a 30 per cent share, and its operating system division has sales in excess of US\$200 million per year, but the company is not without its problems. In 2002 the company made a net loss of some US\$107 million.

SCO News

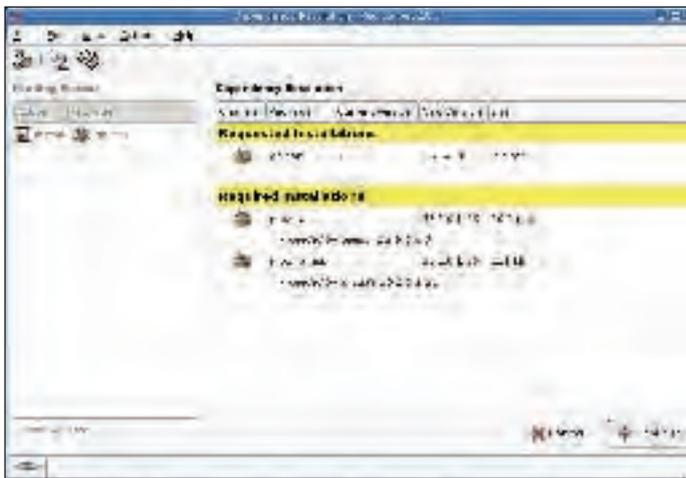
■ Hewlett Packard has been sending mixed messages on the SCO front to its Linux customers. Firstly the company said it would indemnify user of HP systems pre-loaded with Linux against the threat of legal action. SCO rushed out a press release claiming this vindicated its claims of IP theft, something HP denied. Then SCO announced that HP would be the prime sponsoring of its forthcoming whistle-stop US tour.

■ After months of speculation, Darl McBride formally roped SGI into the affair, saying he would revoke the license for the distribution of IRIX. SGI responded with an open letter which said that, in the past, some System V code may have inadvertently made its way into Linux. However, Vice President Rich Altmaier went onto to say that the inclusion of the code fragments was short-lived. "We found better replacements providing the same functionality already available in the Linux kernel. All together, these three small code fragments comprised no more than 200 lines out of the more than one million lines of our overall contributions to Linux. Notably, it appears that most or all of the System V code fragments we found had previously been placed in the public domain, meaning it is very doubtful that the SCO Group has any proprietary claim to these code fragments in any case."

■ Groklaw has continued to highlight the big issues, recently responding to McBride's Open Letter to the Linux Community with an open letter of their own which suggests that SCO executives may, once the case is resolved, find themselves faced with charges of mail fraud and, potentially, pumping stock prices of their own company using misinformation.

■ IBM widened its counter-suit arguing, among other charges, that "SCO made a clear and unambiguous promise to IBM and others that SCO would copy, modify or distribute programs distributed by IBM and others under the GPL only on the terms set out in the GPL."

■ Questions have been raised about the potential for The Canopy Group, a major SCO investor and the ultimate source of the current allegations, to use its position in other Linux-focused companies to mount similar attacks. Canopy has interests in Trolltech, makers of Qt, and Linux Networkx among others.



Ximian's *Red Carpet* software management solution will be included in Novell's Nterprise Linux Services.

.NET ON LINUX

Linux dominates Brainshare Europe

Novell Ximian, instigators of the .NET on Linux project *Mono*, is claiming that its .NET application-enabling software will be ready by the end of this year. Speaking at Novell's European Brainshare conference in Barcelona, Miguel delcaza said "Version 1.0 isn't quite complete, but it will be soon."

Despite the lack of 'official' product availability at present, there is already a quartet of commercial applications – from *Openlink Software*, *Tipic*, *Winfessor* and *SourceGear* – available for the .NET platform, with many more in the pipeline.

Novell also announced that *Ximian Red Carpet* would make it into Nterprise Linux Services 1.0 which is currently in Beta. The package will be

supported on both Red Hat Enterprise Linux and SUSE Linux Enterprise Server providing file, print, messaging, directory and management services in an integrated stack.

"As Novell's Linux strategy rapidly unfolds to customers, the benefits are becoming plain to see," said Chris Stone, Novell vice chairman. "We're moving more of Novell's networking technology to Linux, supporting it along with Linux, and now integrating Ximian products to benefit customers immediately!"

Novell has moved back the launch date of its NetWare for Linux offerings in order to integrate Ximian technology into the package.

For more information on The Mono project, visit www.go-mono.com

METEROLOGY

Watch the Big Blue skies

IBM will be using a Linux-based grid system to monitor conditions in the troposphere. John V. Lombardi, of the US government-funded University of Massachusetts Amherst facility, says the grid will improve weather forecasting and potentially, "save millions of dollars and protect many lives by identifying severe weather systems much sooner than any system currently in use"

The \$40 million facility is the first of its kind able to compensate for the

earth's curvature and obstructions such as large buildings and mountains by using a dense array of short-range radars. These will "communicate with one another and adjust their sensing strategies in direct response to the evolving weather and changing user needs – a dramatic change from current technology!"

IBM says the system has the potential to completely revolutionise meteorology over the next few years.

David Cartwright

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



COMMENT

Legacy systems

“ Sometimes it seems like only yesterday... I started my full-time network manager career supporting a fleet of Sun workstations and servers (and that was before they started calling it Solaris). When a client phoned a few weeks ago and asked me if I'd build them a mail server on a spare Sun that they had lying around, I figured it would be a breeze – after all, I used to do that stuff for a living.

It's fair to say that the installation was no different from any other ground-up Solaris 8 install: stuff the CD in, choose some options, swap discs when prompted for an hour or two, and you're done. It was only when I came to compile and install the mail server (my old favourite *Postfix*, naturally) that I realised how much the Linux vendors spoil us.

Compiling stuff meant finding a downloadable binary of GCC, not to mention the associated libraries, the *make* executable, and so on. At every turn I found myself going back to the Web to download a new bit of stuff that would let me do what I wanted.

All this made me remember my reaction to my first proper Linux install: "Blimey, look at all this stuff that comes as standard – I usually have to download it all myself!" Now I noticed that Sun does ship more stuff than it used to with Solaris (you get some proper shells, for example), and I daresay they've taken it further with Solaris 9, but I still have to take my hat off to the Linux distributors and say: Thanks, guys – I just realised how much grief you save me every time I install this thing! To see what I mean, just try MDK on the coverdisc... ”

Mailserver

Share your opinions, right wrongs and demand justice by writing to *Linux Format*. Reader tips are always welcome – and remember that our forum is always open for discussing your Linux triumphs and woes!

★ Letter of the month

Thanks to the generous folks at Seapine Software, this month's winner receives a full copy of the brand-new **Seapine Surround SCM 2.0**

Linux in schools

Having read Jono Bacon's ideas about Linux in schools (*LXF44 Spreading the Word*), I thought I would share my experiences of trying to implement Linux in my school. I have been using Linux on and off since *Linux Answers* appeared with RH 6.0 on the cover. When I was put in overall charge of IT I thought it would be an ideal opportunity to use my experience and save the school a few quid into the bargain.

The PCs in my school are basically split into 3 groups:

- Networked student PCs and servers;
- Network admin PCs and servers;
- Standalone machines in departments.

I am not actually involved in the teaching of IT and the head of IT has no Linux experience. Therefore, despite my urgings, there was no way he was going to

shift the main fileserver from NT to Linux. So what about the student PCs? Well, school IT is totally reliant on *MS Office – Access* in particular – which means even the use of *OpenOffice.org* is pretty much out of the question.

I thought perhaps I would have more luck of the machines I have direct control over.. Well, in short, I didn't. Most of the admin staff are still unhappy with the switch to *MS Office from Word Perfect*, so now introducing Linux would throw them into a complete spin. The fileserver then, surely that could run Linux and use *Samba*. Nope. Unfortunately it also has to run *SIMS (School Information Management System)*, which only runs on Windows and has recently begun to use *MSSQL*, so that is a no-go area as well.

By now I was getting desperate. What about the stand-alone machines? They are mainly used for office tasks, so Linux might

have a chance. Again I was foiled as they are also used to run proprietary applications specific to individual subjects which (not surprisingly) run on Windows. I did think of Wine but the users are not really up to it, even if all the programs would work. I did, however, score a minor victory by installing *OpenOffice.org* on all machines that do not need access to *Access*.

I had just about given up when I found that I needed two firewall/routers and a content filter for various parts of the network. Enter *IPCop with DansGuardian* added in for good measure. I hijacked two old PCs, installed *IPCop* and within a few hours we had two fully functioning firewall boxes. At last – some demonstrable success for my cunning Linux plan!

To sum up, I am all in favour of using Linux in schools and have tried very hard to use it pretty much everywhere. However, until

there is wider acceptance of Linux, less reliance on Microsoft products in examination subjects and a greater level of training, it is going to remain very difficult to implement Linux on a wide scale in UK schools. Of course, I have no intention of letting the head of IT buying Windows Server 2003 when he wants to upgrade, maybe SUSE would do the trick...

Frank Johnson, *via email*

Thanks for your inspiring, if not entirely successful story. It's a shame that areas which could make better use of money than spending it on software licences are left with little choice. The only way to improve the situation is through, perhaps ironically, education. This is one area where many European nations are ahead of the UK, due to positive initiatives by the government in countries such as France and Germany.

As a reward, please accept a copy of *Seapine Surround SCM*, worth a whopping US\$595.

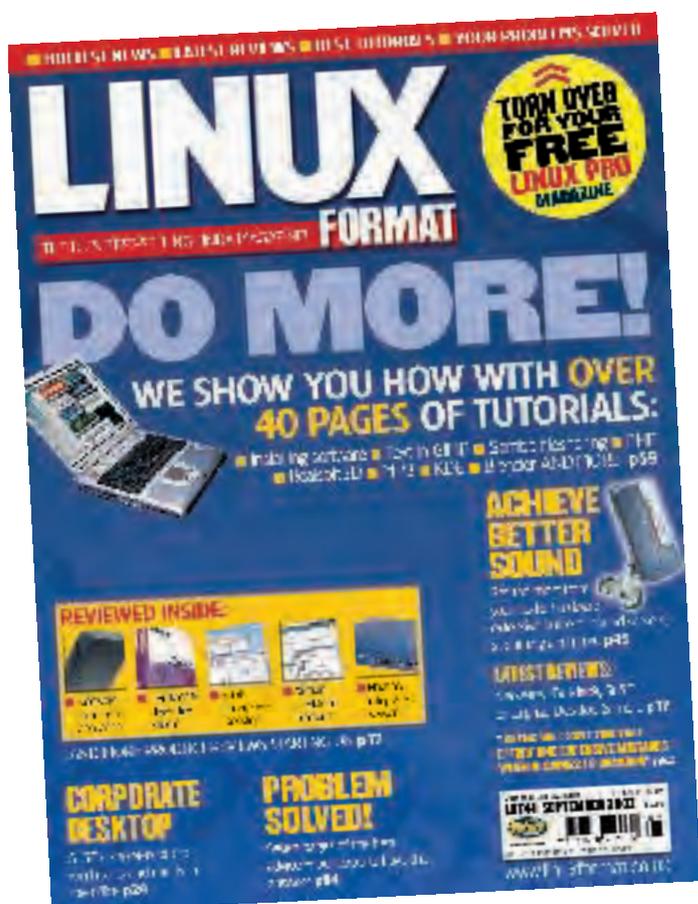
POW! It's Evesham

About 18 months ago I stuck my neck out and convinced the management of the small company I work for (65 people based in six Centres around Powys, we're a training organisation attached to the local County Council) to replace our NT4 and elderly NetWare servers with Linux-based machines. I knew nothing much about Linux other than what I'd read and, due to the fact that I've only had time to learn the things I need for work, I still consider myself very much a Linux 'newbie', even

though I hate that term! Recently I decided to get a machine at home to run Linux on, so that I can get familiar with the various desktop



Evesham not only sells Linux desktops, but also looks after its customers.



options and different flavours of Linux (I want to get some of our training rooms dual booting Windows and Linux to be able to offer Linux 'taster' days to local people and businesses).

I just happened to buy *LXF43*, I've been picking up the odd copy for a while though they're hard to come by in mid-Wales, and read the article reviewing the Evesham Mandrake-based PC. It was impressive reading and I thought, what the hell I could do a lot worse than this machine.

I placed an order with Evesham and the kit arrived a few days later. The specification of the speakers has been upped and now comes with 4 satellites and the bass speaker. The machine functioned fine out of the box and I was impressed by the quality and speed. It also came with a brief but very satisfactory re-installation guide for the MDK set-up.

After a couple of days I tried loading a CD and ran across a problem. I couldn't access it. On investigation I couldn't access DVDs either and on further investigation I discovered that the drive wasn't being recognised. I did a visual check to see if the cables looked as

if they were connected, which they were, and then decided to let Evesham sort it out as it was only a few days old.

I telephoned Evesham on a Wednesday and explained the problem to a cheerful customer support person who told me to hang on for a few moments. Within half a minute, he was back to tell me that an engineer would be with me on Friday. "I have to tell you I live on a small-holding up a track in the middle of nowhere a couple of miles from Builth Wells" I said. "We'll get the engineer to give you a ring when he's nearly there so you can guide him in" he replied. Two days later the engineer rang from Builth, found the place without too much trouble and replaced the drive. He was good natured, friendly and knowledgeable.

So, thanks to Evesham for a first-rate service and a big thank you to *Linux Format* for recommending them so accurately. I've since taken out a subscription for *LXF* so I don't miss any more issues and I'll be using Evesham again.

Ray, *via email*

Blimey, the Christmas spirit has arrived early – usually magazines are full of

We cleared up a few myths in *LXF44*, but didn't investigate their source...

letters complaining about companies, not praising them. Glad to hear the Evesham has lived up to expectations.

Call my Bluff!

RE: *BLUFF!* (*LXF44*, page 52) Fewer files in your directories is better...

Back in the bad old days, over 20 years ago, when people were running V.7 or even SYS III, before the Berkley Fast File System got ported into every manufactures 'nix, we had the AT&T filesystem. A directory entry consisted of 14 bytes for the files name and 2 bytes for a pointer to the inode that pointed at the blocks in the file. Yes, we had really small disks.

Since this structure didn't allow for many entries in a directory, or for that matter many blocks in a file, there were first, second and third level "extends" which were blocks of inodes that pointed at blocks of inodes which... You get the idea. Not very efficient.

If you had a directory that had more than 1024 files in it, you where into second or third level extends. So your OS had to read the block (the directory) to determine the first level extend and read that, to determine the second level and read that, to determine the third level extend and read that to find the block that contained the file name you were looking for.

In days of old before BSD's fast files system came into common use, having a large number of files in a directory did slow down the system. Jim Kissel, *via email*

You are quite correct that this used to be a problem, and is probably why the myth still persists to this day. Thanks for the explanation!

Is Windows UNIX?!

Though I am a committed Linux user (SUSE), I have over the last few years installed, reinstalled, upgraded and repaired NT, 2000, XP a great many times on the machines of others. Watching the files go in, observing the functions, processes and relationships between them (insofar as that is possible beneath the heavyweight encryption and a rather cumbersome GUI), I have wondered many times just what lies beneath the surface of Windows.

SUBMISSION ADVICE

WHAT WE WANT:

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

WHAT WE DON'T WANT:

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

WRITE TO US AT:

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

Am I entirely alone in thinking, as I have done often on these occasions, that what I am actually looking at is a cut-down Unix/Linux base with a not very well crafted kernel, all cloaked by industrial strength encryption? Perhaps I am quite wrong in this, but the similarities appear to extend to more than mere coincidence, or the requirements of a common hardware architecture.

Apropos all of this, I have been watching with considerable interest these last few months, the extraordinary, almost laughable, activities of SCO, apparently backed by Microsoft. Like everyone else, I automatically assumed that this was yet another attack on Linux which is beginning to pose an increasing risk to Redmond's long-term viability.

However, if there is, in fact, a reworked Unix/Linux foundation to the NT family base, this puts an entirely different complexion on the whole matter. In these circumstances, those curious activities by SCO become not so much an attack on Linux as the beginnings of a legal and commercial defence of Windows composition in order to ultimately strengthen its own monopoly... which is an interesting thought!

Allan Lewis, *via email*

It is no secret that large parts of older versions of Microsoft OSes were built with code from BSD (Not Linux, being definitely a flavour of Unix), including most notably the TCP/IP stack and attendant programs such as *ftp*, *finger*, *nslookup* etc. Under the terms of the BSD licence this is perfectly permissible, though it does obviously make you wonder about various anti-free software speeches given by Microsoft bigwigs. >>

READER TIPS

JOE TIPS

After having read *LXF44*, I wanted to mention some important details were left out in regards to the *Joe* editor. I think the main reason the *Joe* editor has any following at all is because it is so familiar to former *WordStar* users. Back in the CP/M and early DOS days, *WordStar* was quite widely used, and the *Joe* editor is like slipping on an old comfy set of slippers because of that.

There is also the ability to emulate several other editors. With the one binary, just symlinking it to a different name and running the editor from that link will turn on the wanted emulation. Here's a list:

- *joe* – Standard *Joe*, similar to *WordStar*.
 - *jstar* – Much closer *WordStar* emulation, with enhancements.
 - *jpico* – Just like *Pico* (*Pine's* editor) and *Nano*, with enhancements.
 - *jmacs* – GNU-Emacs emulation.
 - *sjoe* – Restricted *Joe*, limits editing to files specified on command line, handy for secured scripts that are needing an editor.
- Lastly, *Joe* lets you edit multiple

files simultaneously by popping between text windows. Handy for copy and paste between files.

Brandon Darbro, *Seattle, aka TheDarb* of www.PCLinuxOnline.com

MDK CONTRIBUTING

As a 24/7 Mandrake user, I read with interest Jono Bacon's comments on contributing to Open Source projects. I would like to comment on some of the ways in which we, the users, are making our contribution.

Jono did, of course, mention Mandrake Club. Many would agree with him that the need to request money in this way is unfortunate. However, it is worth noting that MandrakeSoft has always believed that its full distro should be absolutely free in cash terms as well as in Open Source terms. Boxed sets are of course available, but every bit of the distro can be downloaded as soon as it's released. For a small company to maintain such high standards, revenue has to be obtained from elsewhere. Voluntary contributions are a legitimate source of income, and many of us support it in this way because we believe

that Mandrake could not afford to continue such largesse unless we do give monetary support.

That apart, there are many other ways of supporting. As Jono pointed out, documentation is often woefully inadequate, not least due to the fact that the hardware situation changes daily. The only way that this can be addressed is by users contributing time and effort. Mandrake did attempt to start a hardware compatibility list, but found the task impossible. What was achieved remains available, as mentioned in your reply to T.W. Groves (*Mailserver, LXF45*). There is however, an incomplete but much more up-to-date list growing on the users' TWiki site (yes, Jono, we have already started to address this!)

The TWiki pages were begun around six months ago, growing slowly at first, but recently the growth has been fast, so that we have reached the situation where the location of the information is not always self-evident. We are, therefore, in the middle of reorganising the index to address this. The original URL of the TWiki web was difficult to remember, so

one member has donated a domain to redirect you to the introductory page. Anyone who uses Mandrake, would benefit from browsing our site <http://twiki.mdklinuxfaq.org>

Finally, I would like to remind users that this is not a closed 'club'. We welcome material from any user, particularly if they are able to add to the hardware compatibility pages, or able to write a mini HOW-TO on any of the frequently encountered installation problems. Anne Wilson, *via email*
See page 108 for an exclusive *LXF* money-off offer for Mandrake Club!

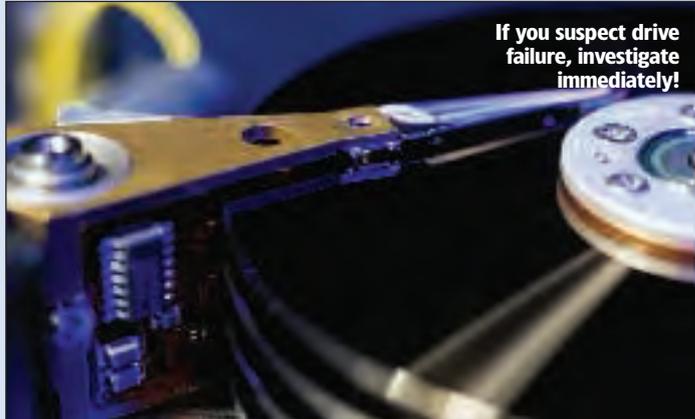
CLUNKING DRIVES

LXF42 page 96 – freezing after hearing a clunk from the drive. One thing that will be worth checking is the power supply. I've had the same problem 3 times in the last 12 months and all of them have due to poor voltage regulation on the +12 volts. The first time this occurred, the system was not behaving itself for a while; when I heard the drives park their heads, followed some seconds later by the sound of a number of drives restarting, with the system then freezing up.

Helpdex

shane_collinge@yahoo.com





Because in my case more than one drive was affected, I checked the voltages via *lm-sensors*. This showed the +12v down slightly. On dismantling the system and monitoring the voltage with a voltmeter showed the voltage to be around 11 volts, and when the voltage dropped to about 10.7 volts the drives would park the heads and switch off. Replacing the power supplies fixed all three. All three times, once the drives had switched themselves off, the file system would not restart when the drives had restarted. The only option I had was to reset via the case or power off.

Lindsay, Victoria, Australia

Power supplies don't last forever, especially desktop ones subjected to high demand. Well done for spotting a problem which many people overlook. However, you should also check that the PSU is up to the job (*ie* has enough Watts coming out), as this is the usual cause of dropping voltages.

KILL SPAM DEAD!

Currently, many spam emails use HTML pages or tables for their content because they can then break their words up and prevent the operation of word filters. In addition, an HTML message is much longer compared with a text email, so a series of messages that might normally be a 15-second download now takes two minutes.

Recently I shifted from *Evolution* to *KMail*. Interestingly, I found that *KMail*'s default mode of opening an email shows the code, not the web page display. It was this which suggested the potential filtering

method because I suddenly realised that virtually everyone who is using email for normal communications (or at least those with whom I correspond by email) uses plain text – only the spam merchants are using HTML and Tables.

As a result, there is an incredibly simple way to beat them at their own game and filter off a very large percentage of their garbage into a spam folder for destruction. All you do is set up filters that specifically look for HTML codes in the body of a message. For example, currently my filters are set to look for <HTML>, <html>, <TABLE> and similar codes that indicate a web-style page is coming in. So far, the filters are having hit rates that are very high – between 70 and 80 per cent and I am sure I can improve this as I set more HTML filters. They also appear to be very fast. And of course, if the spam merchants think that to beat you, they will revert to plain text, and thus normal word filters will work once more. You can't lose!

As a further assisting item, notify all your email contacts that from this point onwards, you are forced to treat all HTML messages as potential spam and that you prefer **ONLY** plain text messages unless the sender has previously notified you that there will be either an attachment or an HTML encoded message. You may not be able to stop receiving spam, but at least now you have a darn good chance of easily sending it to the garbage folder.

Dr Tony Young, Queensland, Australia



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

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

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

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



WWW.SUSE.CO.UK

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...

Xinit Sharqstation >>

Serial ATA on the desktop! A well-respected server manufacturer dips its first toe into workstation waters **p24**

MT Studio

Can the performance of this rapid application development environment measure up to its initial conceptual promise? **p27**

Lazarus

Aiming for similar functionality to Delphi or Kylix, this looks loaded with potential prior to its 1.0 release. Personalised update flags endeared it to us immediately! **p30**

SUSE 9.0

MIDI capabilities, *Scribus* DTP, *Kopete* IM client, *OpenOffice 1.1*, and User-Mode Linux support are just a small selection of the new components **p32**



Linux peripherals

If you play games on Linux and are serious about getting high scores, your fingers will be itching to use force-feedback joysticks, steering wheels and joypads... **p34**

Books

Extreme Programming With Ant, *Beginning Red Hat Linux 9*, *A Practical Guide to CURL*, *Mastering Red Hat Linux 9*, *Essential CVS* **p36**

LINUX FORMAT BENCHMARKS EXPLAINED

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

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

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

BENCHMARKS

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

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

LINUX WORKSTATION

Xinit SharqStation

Xinit's a pretty big fish in the server ocean – can it extend its fearsome reputation into the workstation pond as well? Paul Hudson tests the water...

BUYER INFO

A high-power server with lots of hard drive performance. Also consider Armari's Pro3D Workstation, reviewed in *LXF46*

- **SUPPLIER** Xinit
- **PRICE** £3485+VAT
- **WEB** www.xinit.com

The two machines we've seen from Xinit Systems in the past have both been high-performance, disk-oriented servers that have performed excellently in our tests. Along with the company's recent attack on the NAS market (see *Linux Pro*, page 14), Xinit is also moving with force into the Linux workstation market – still one of the least tapped arenas around, despite adoption of our favourite OS by major players in the animation market, like ILM and DreamWorks for instance.

Looking over the specs of this machine, it's clear that Xinit is yet again pulling out all the stops to try to make a splash in the biggest possible way. Whether or not their commitment to the Xeon architecture will help or hinder their plans will be seen!

Back in the day

Workstations in Linux are still somewhat of a rarity, with the market still trying to decide quite what is the best route to take. Our first Linux workstation, reviewed just last issue, plumped for the Opteron CPU – a choice that reflected very well on its CPU performance. This machine, however, goes with a much more traditional Xeon solution – we've found Xeons to be much slower than Opterons in our tests, but it's possible this choice may still be borne out by the quality of the other hardware in the box.

Backing up the twin 3.06GHz CPUs is 2GB of ECC RAM, four 36GB SATA hard drives, and two gigabit

Ethernet adapters. SATA is really starting to catch on now for desktop machines, and this model comes preconfigured for RAID 0 (disk striping without parity) for maximum performance, all connected up via a 3Ware SATA RAID controller.

On the multimedia front, the system comes with a 128MB nVIDIA Quadro4 980XGL – the top-end graphics card from nVIDIA's Quadro4 range, which runs at AGP 8x. While not as powerful as the QuadroFX we looked at last month, this is still more than enough to handle nearly everyone's graphical requirements. There's also a Soundblaster Audigy card in there which is a new sight for us, as well as a DVD-ROM and a CD-RW to boot.

What this machine misses is a high-quality monitor to back up the graphics card, and also note that you'll need to find your own peripherals. Clearly the high point of this machine is its hard drives – four SATA drives hooked up through RAID 0 should provide excellent performance, particularly because these are 10,000RPM models. The question is, will this balance the potentially poor CPU performance?

Ooh la-la!

Eager to put the SharqStation through its paces, we set it all up and powered on. The fans themselves aren't too obtrusive, although they are by no means quiet; but the main noise source was a peculiar buzzing that fluctuated in volume seemingly randomly. At one point the vibration was so loud that we had complaints coming in from people over twenty metres away – it literally did sound like a fan somewhere was sawing through the case! When we spoke to Xinit about this problem, we were assured that it was merely a result of the



Is black the new beige? Some might say that the workstation market in general has some catching up to do in terms of case design – even a simple shroud to protect the drives from dust would be nice...

machine being an early test, and that all machines sold to customers would be much quieter and vibration-free.

The machine booted very quickly, because, as we later discovered, the boot process had been streamlined so that only essential services were loaded – a boon for security as well as speed. Further rooting around showed that the Red Hat 9 installation was very well thought-out, and we actually had to do no tweaking or installing to run our benchmarks. While no CDs

came with the box, Xinit had at least taken the time to make sure that much of the most popular software was all installed and ready for use.

On the performance front, there were a few highs and lows. In the CPU-based test, *ogenc*, the machine returned a respectable enough score, but certainly nothing stunning – the 2GHz Opterons we review last month managed a score 50 per cent faster, which is in line with our expectations. Of the rest of the scores, only really the



As well as a couple of USB ports on the front of the case, there's a couple round the back too. Along with the usual selection of connectivity, two gigabit Ethernet ports promise nimble networking capabilities.

hard drive score is of interest – while 2.19 is not stunning by any means, it at least inches ahead of the pack somewhat. The big disappointment is the MySQL test – 0.63 is very poor, and it returned approximately the same no matter how many times we ran the test. While it's possible we could eke more performance out of MySQL if we tweaked various bits and pieces, it's not our testing policy to attempt any optimisation beyond how the machine was shipped, as this might bias results.

These are average results to say the least, and quite disappointing. We expected a lot more from the hard drives, and clearly there's a problem with the MySQL/PHP install somewhere. While the raw CPU performance shows that this machine has a lot of potential, the other poor results drag the overall score back to a fairly time 2.46. Lacking stellar performance, this machine is left looking for a niche it can fit into. One of the high-points of

the system (and not something we actually test in our benchmarks simply because of the immense variability of the machines we review), is the graphics card. The Quadro4 in this box is still an excellent card despite its age, and this shines through nicely because Xinit took the time to properly configure the graphics to work with the card – the nVIDIA Linux drivers are installed and working perfectly. To make this point clear, the machine came pre-installed with *Celestia*, a 3D galaxy rendering system, *Blender*, everyone's favourite GPL 3D tool, and *Enemy Territory*, the first-person shooter game. All three ran like a dream, particularly *Enemy Territory*, which we spent many long hours testing to our hearts' content, we can assure you!

Added to the top-end graphics is an excellent sound card that comes configured ready for use. This will please gamers no end, we're sure, because the combination of smooth video and crystal-clear audio means that this is a machine you can buy and will work from the off with your games. One minor drawback is the selection of installed packages – they might seem a little random at first, because programs like *KDevelop* are in there whereas, despite the presence of a DVD-ROM, no DVD player program is there. Hopefully this will be sorted out in the final release, as it shows a little packaging bias on behalf of the designers!

El Capitan

While it's true that the graphics are indeed top-notch, few people are going to use this machine for gaming. Instead, the super-fast video rendering will be of most importance to 3D artists who don't have time to sit around waiting for things to draw. To back this up, the relatively speedy disks means that data transfer won't be a problem for moderate loads – this is probably where the machine fits in best.

However, the competition is fairly tight in the arena. For graphics performance, the Armari box we reviewed last issue comes in £100 cheaper, and yet manages to pack into there a better graphics card and much faster CPUs, as well as a high-quality monitor and, perhaps critically, an enterprise-level OS (SLES 8). The downside of the Armari machine is that it doesn't have SCSI or even SATA hard drives, which means the data transfer isn't as fast. On the basis of

pure data transfer, the Fujitsu server we reviewed last issue has much better data transfer, equally as fast CPUs, more RAM, and substantially faster hard disks.

By the time this box gets released to clients, we hope that Xinit will have been over it with a fine-toothed comb to make sure the performance glitches we found are eliminated entirely. While this machine is never going to be a world-class performer because it's held back by the Xeon architecture, a lot more can be done to maximise what's there already.

On the money

For this price, you'd be forgiven for questioning how much value you're getting. However, this is really the only machine currently available that combines digital content creation abilities with hard drives powerful enough to give the performance that's demanded in today's environment. Once second-generation SATA comes out, Xinit's decision to use the more expensive technology will be borne through as their machine will be able to scale up easily. Having said that, while the hard drive performance would be nice, users will undoubtedly see a great more performance if Xinit were to take the leap to AMD64.

Overall, this is quite an expensive machine, owing largely to the high-end hard drives bundled as standard. The question remains whether you need that extra disk transfer performance – if so, this machine is as good as it gets right now. [LXF](#)

BENCHMARKS

HD	2.19
AB	3.33
MYSQL	0.63
COMPILE	1.68
OGGENC	4.45
OVERALL	2.46

LINUX FORMAT VERDICT

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

A Xeon workstation with server-level hard drives should keep artists happy, but the price tag might make you want to shop around first.

RATING **8/10**



WEB-BASED RAD

MT Studio Standard 'Plug-and-Play' Bundle

Richard Drummond has used quite a few IDEs in his time, so how does this compare to the established standards?

BUYER INFO

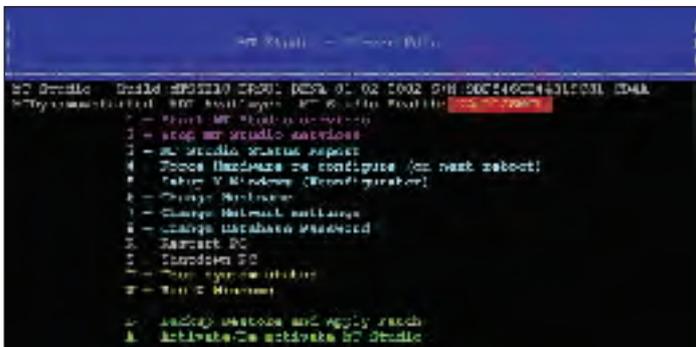
Competing solutions for creating web-based apps with rich interfaces include *Omnis Studio* and *Cold Fusion Studio*.

- **PUBLISHER** Mart Technology Ltd
- **WEB** www.marttech.com
- **PHONE** +44 (0) 20 8429 7345
- **PRICE** £595.00+VAT
(Single developer licence)

Rapid development of web-based applications is not a new concept, but Mart Technology thinks that it has a new spin to boost developer productivity and underline the 'rapid' in RAD. Their Java and XML-based *MT Studio* suite is a complete package which includes an application server and integrated development environment. The novelty is that the IDE is itself a web application served up by the very same server used to deploy the applications that you use it to develop. Not only does this emphasise the power of *MT Studio's* technology (it's sophisticated enough to host a complex IDE), it also cuts out the deployment stage of application

development (you develop applications on the server) – and can potentially offer benefits for team development (though we tested a single-user version).

Hang on a minute, I know you're thinking, a web-based IDE? An IDE using HTML forms – that can't be much fun. Well, think again! While *MT Studio* can deliver conventional web-based applications, that is, pages that employ static or dynamic HTML (of course, optionally augmented with server-side scripting, CSS, JavaScript and so on), the story doesn't stop there. It can also serve up complex graphical interfaces that are created at run-time in the client's browser via a Java2 applet called *MTClient*. (*MTClient* requires a browser equipped with the J2SE 1.4 plug-in, but no other client-side installation is required.) Thus your web applications can employ the rich set of widgets offered by Java's *Swing* toolkit, rather than the feeble GUI controls that HTML allows. It is *MTClient* which is used to present the *MT Studio* IDE to the developer.



The 'Plug-and-play' Bundle boots up to a management console for doing basic administration of your *MT Studio* server.



MTWizard can even be used by somebody with no real programming experience to develop simple classes of application.

Full environment

MT Studio Standard is offered in two guises: as a suite which can be deployed on Windows, Linux or other Unix-like systems; or, the version we tested, as a 'Plug-and-play' Bundle. This latter is a neat concept. It includes a live Linux CD (based on Red Hat 7.3), *MT Studio* itself and Sybase's *SQL Anywhere* database. It requires no installation as such, and you can use it to boot up a full *MT Studio* environment on any PC – no matter what operating system you have installed on it. It does require disk space for storing applications and settings, but it can use any Linux ext2 or Window FAT or NTFS partition with sufficient space for this (actually, NTFS partitions will require some prepping before use).

MT Studio Standard consists of four core components. Firstly, *MTDynamo* is the heart of the application server – and is actually implemented as a servlet hosted by the *Apache Tomcat* servlet engine (although it should work with any other J2EE compliant server). *MTDynamo* talks to a database back-end where your application data and pages are stored and serves requested pages to the user's browser. *MTClient*, as we mentioned above, is the client-side part of the equation and is needed to create GUI-based forms at run-time. *MTIDE* is the development environment proper, itself an application delivered via *MTClient*. Finally, *MTDesigner* is a GUI-based designer for creating the graphical forms which are deployed with *MTClient*. As a bonus, a tool for template-based development, *MTWizard*, is also included. With an appropriately prepared template,

While the 'Plug-and-play' Bundle sounds like a good idea, it was where my problems with *MT Studio* began. When you boot the CD for the first time, it looks for a suitable partition for free space to use. However, you don't get to choose – it automatically uses the first one it finds. In my case, this happened to my ext3 /usr partition rather than the ext2 partition that I



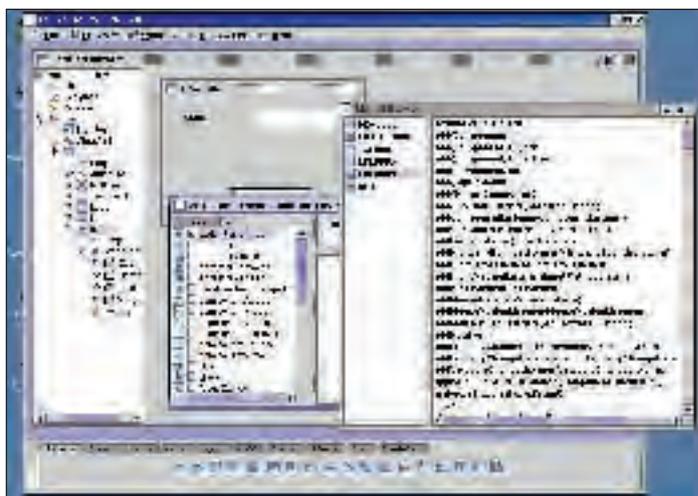
REVIEWS MT Studio Standard

« created especially for *MT Studio* to use. It duly mounted my *ext3* filesystem as *ext2*, did its stuff, and in the process rendered the bootloader on that partition inoperable. Although easy to use, by its nature the *'Plug-and-play' Bundle* is considerably less flexible than the conventional version.

When the *MT Studio* live CD boots up, you'll see a text console displaying a menu-based management program. From here you can start and stop the server, print various diagnostics, perform various reconfiguration tasks, launch X, and so on. The live CD contains a full X distribution, with *Mozilla* and dozens of other applications. You can use this to launch a browser and develop on the same host as the server, but I don't recommend it. Running from a CD is slow anyway, but running both server, X and a browser-based client from CD on the one box is really slow – unless your system boasts an obscenely large amount of memory.

Development

To start developing with *MT Studio*, you simply point your browser at the IP address of the server, and up pops the home page with links to the various components. *MTIDE* will be the focus of your work with *MT Studio*, and this provides tools for managing and editing pages, the basic building block of applications. A page can be a garden variety HTML page – embellished with client-side JavaScript or server-side tags and scripting – or an XML page – which, since a GUI-based form is embedded in a page as XML, is required for creating GUIs to be deployed via *MTClient*.

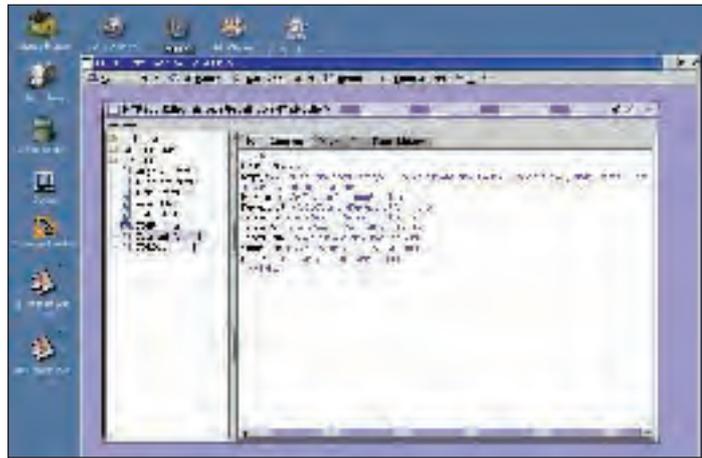


MTDesigner is a fairly sophisticated form builder for creating GUIs to be displayed via *MTClient* in a user's browser.

Serving-side scripting is effected in *MT Studio* with *MT Tags*, extensions to HTML similar to *Cold Fusion's* CFML or JSP tags. These allows you to access form or sessions variables, do simple operations and conditionals, embed SQL queries and format SQL result sets as HTML. Also, the *MTSCRIPT* tag allows you to include arbitrary logic in a page, written in *BeanShell*, a scripted variant of Java (see www.beanshell.org). *BeanShell* is syntactically identical to Java, so the potential for code re-use from existing server-side projects is high.

MTIDE provides controls for creating and manipulating pages and for simple project management. It includes a page editor that can colourise HTML tags and *MT Tags*, but is otherwise very primitive. Annoyingly, it only lets you edit one page at a time, and if you accidentally switch to another page without saving, it will forget your changes. It also provides no function for previewing your pages. With the single developer version, you must disconnect from the server (only one user may connect to the server at a time) and open a browser to window to view the fruits of your labours. The editor isn't terribly quick anyway; this just slows the whole thing down further.

Java-based GUIs are created with *MT Designer*, which presents a friendly environment for building GUI forms. If you've used tools such as *NetBeans* or *JBuilder* before, *MTDesigner* will be instantly familiar. It's obviously not as powerful as those full-blown Java IDEs, but for this job it's just as usable. It provides the usual palette of components with you can 'paint' on your form to construct your interface



This is *MTIDE* with the page editor open. Notice the *MT Tags* used here for accessing the web server's environment.

and a property editor to tweak component settings. This also includes tools for constructing and implementing event handler to breathe life into your GUIs. Event handlers are also written in *BeanShell*, so no retraining for Java developers is required and, similarly, you can potentially re-use code from existing client-side code. However, with the single developer license you cannot launch *MTIDE* and *MTDesigner* simultaneously, which is a nuisance because *MTDesigner* embeds the GUIs it creates as XML in an XML page. Clicking on such an object in a page should launch *MTDesigner* to edit it, but it doesn't.

Conceptually, *MT Studio* could provide a very fluid and fast environment for building web-based apps. Unfortunately the version I tested was anything but. The primitiveness of the editor, the inability to run *MTIDE* and *MTDesigner* at the same time or to be able to easily test the pages you are building means that work-flow is severely restricted. Combined with a general lack of stability – seemingly random Java exceptions would bring the system to its knees, reporting only a cryptic error code – and you're left with a package that is nowhere near as user-friendly as it should be.

Another crucial problem with *MT Studio* is lack of documentation. A useful printed install guide and user manual are included, but while the latter covers the architecture and general use of *MT Studio* well, it lacks any meat for doing real application development. *MTIDE* offers a help menu, which can launch a Java help viewer – but alas with no content – or open a browser window on the developer section of the Mart Tech

website – which contains links to non-existent documents. The plan is obviously at some stage to provide a reference manual, but currently without one it is difficult to make good use of this product. No sample applications beyond the trivial are included either.

Thus, it's difficult to fairly rate *MT Studio* as a product. On paper it looks great, and the architecture is ingenious; but I just wasn't able to satisfactorily use it – partly due to problems discussed above and partly due to a lack of docs. As a solution for developing dynamic HTML applications, even if it worked perfectly, I don't think *MT Studio* has a great deal to offer over the competition. But for developing 'thicker' GUI-based applications that can be served up via a browser, Mart Tech could really be on to a winner. As such, and with its licensing model, it's not suitable for developing public websites, but for Intranet work and for all kinds of in-house bespoke programming where you need to deliver a database-driven application quickly to a diverse set of users, *MT Studio* could really be attractive and potentially head-and-shoulders above rival products – if only it worked better! [LXF](#)

LINUX FORMAT VERDICT

FEATURES	7/10
PERFORMANCE	5/10
EASE OF USE	7/10
VALUE FOR MONEY	7/10

MT Studio is a conceptually exciting product, but unfortunately I experienced too many problems with this 'live' version to be able to recommend it. Your mileage may vary though...

RATING **6/10**



NEW DISTRO

SUSE Linux 9.0 Professional

It's been six months since we last looked at SUSE's leading desktop distro, and, on schedule, the next version is here. Paul Hudson finds out whether it's worth the upgrade...

BUYER INFO

Powerful desktop distro with lots to offer everyone. Also consider Mandrake 9.2 or Fedora.

- **SUPPLIER** SUSE
- **PRICE** £60
- **WEB** www.suse.co.uk

As desktop distros go, SUSE 8.2 pretty much had it sorted – it was easy to install, packed with functionality, and easily the most comprehensively documented distribution that was around. Given all this, it's unsurprising that we gave it our coveted *Top Stuff* award, as it easily outclassed its competitors. In our exclusive feature last issue we travelled to the SUSE HQ in Germany to get a behind-the-scenes look at the development of the next iteration of SUSE Linux, and saw an exciting product just waiting for a chance to prove itself. Now, at last, SUSE 9.0 is available, and we have the first copy...

Software roundup

While there are lots of changes to SUSE 9.0 that aren't available elsewhere, it's clear that the majority of changes come

from the fact that it includes all-new versions of just about every package that's bundled. The headline package is, unsurprisingly, *OpenOffice.org 1.1*, which offers a raft of new features over its predecessor. If you're scratching your head and wondering whether SUSE have invented a time machine, yes, it is true that *OpenOffice.org* only came out in early October and, due to production deadlines, SUSE had their gold CDs for SUSE 9 available in the middle of September, which means that the final release of *OOo* may differ from the version SUSE 9 comes with. Fortunately, the fixes in the final code for *OpenOffice.org* are small, and are unlikely to affect anyone.

Also new in SUSE 9 is the soon-to-be-ubiquitous DTP program *Scribus*, which attracted a solid review in LXF45 and is likely to be a common feature in distros as new versions are released. The kernel has now been updated to 2.4.21 (with SUSE's usual 2.6 backports), KDE 3.1.4, and a few new faces too – *Kopete*, the KDE instant messenger client, is now included as standard, as is User-Mode Linux support for more advanced users.

Perhaps the biggest features in SUSE 9.0 are the ones that most people won't even see – in fact, we



probably wouldn't have noticed them if we didn't get a full tour around the features when we visited SUSE in August. For example, SUSE 9.0 includes an advanced set of programs and features geared towards musicians looking to utilise their machine as a recording and mixing station for their MIDI kit – while this is all Open Source, it was coded largely inside SUSE and this is its first airing in the community.

As with SUSE 8.2, 9.0 comes with five CDs and one DVD packed with programs ready for installation through *YaST*. Judging by the number of disc changes we had to do when performing a basic install it seems that SUSE have gone to great lengths to use up as much of the available CD space as possible, and it's likely that one of the next few releases of SUSE will require an extra CD for all the new packages.

Installation woes

Despite so much having gone into quality assurance, we still had several problems getting it to work on our

machines. Three machines, all of which are of standard build and had run Debian, Mandrake, and Red Hat in the past, failed to get past an early installation screen with SUSE 9. The last spare machine we had in the office finally made it past the installation, and succeeded in booting up into KDE.

Once things were working, we had time to look around the default install and were pleased to find that it's as comprehensive as ever – everything from *Qt Designer* to *The GIMP* was installed, along with tools that have sometimes been mysteriously absent from SUSE builds such as *flex*, *bison*, and *lynx*. Despite all these packages being installed by default, precious few services run at startup – only really the SSH daemon is there, which is pretty much standard.

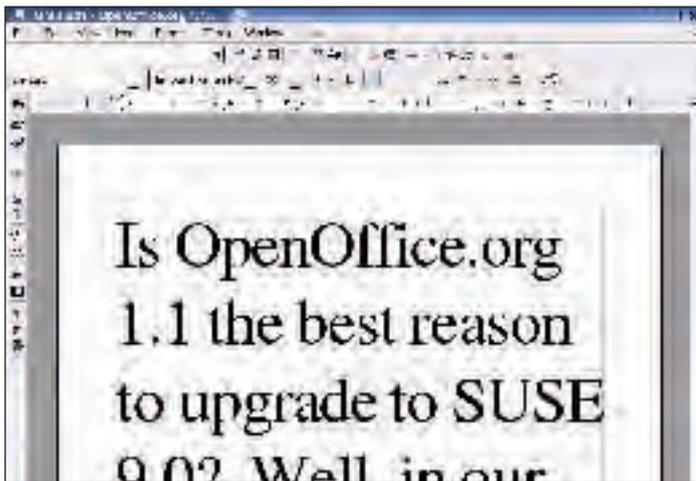
We imagine the installation should be a great deal easier for most people, though, particularly now that so much work has gone into making the distro support more WinModems as well as



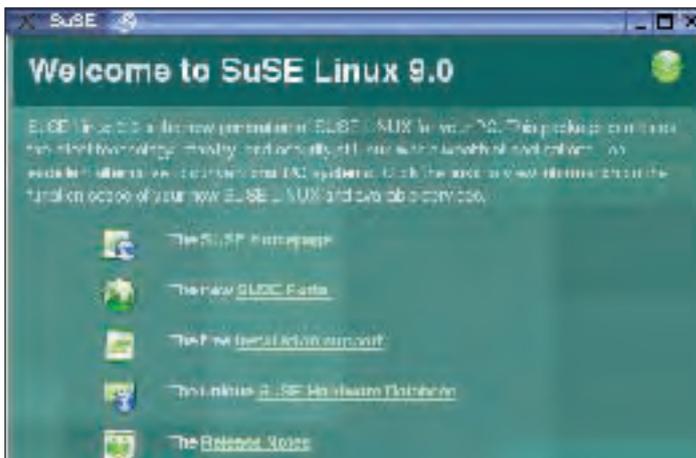
The *YaST* interface is attractive and getting better all the time – if only there were fewer acronyms!



KDE 3.1.4 running on SUSE 9.0. Yes, it looks like any version of KDE 3.x running on any distro, but hey – it's only £50, right?



The headline package in SUSE 9.0 is OpenOffice.org 1.1 – an excellent upgrade to a top office suite, and should definitely help shift some boxes.



The new SUSE welcome screen looks pretty enough, but it mixes up SUSE and SuSE enough to give you a headache!

ADSL connections. Windows users looking to dual-boot SUSE will find NTFS resizing has now been implemented fully, which should hopefully eliminate one more barrier to installation.

Tuning and tweaks

Above and beyond the latest upgrades to the software bundled with SUSE, there are a number of fixes, tweaks,

and upgrades to the SUSE-specific code behind the distro. The key changes are visible in *YaST*, where various new modules have been added to ease administration. The main interface to *YaST* is very Mandrake-like now, which is good to see as Mandrake's design has always been very easy on the eye. However, the individual module screens are as grey and unappealing as ever, which

make them look less friendly and easy to use than they actually are.

Configuring *Samba*, NTP, DHCP, and other complex servers is now easy thanks to the improved layout of *YaST*, and even home network admins should be able to get a powerful server up and running based entirely on the manual bundled in the box. Having said that, the manual is still a requirement because it's easy to drown under the wave of acronyms that *YaST* relentlessly throws out without explanation.

Real-world changes

Even once you take into account the big changes in the packages as well as the *YaST* changes, there's even more to SUSE 9.0 than first meets the eye. Perhaps the most interesting change for most users will be the addition of a snapshot of the 2.6 kernel readily available for more adventurous users. While it might not be stable by any means, at least it gives people the chance to get a glimpse of the future – programmers can make sure their code is compatible, and everyone else can just tinker with the latest and greatest Linux has to offer.

Another big change for SUSE is the availability of an Athlon 64-enhanced port from launch. When Opteron was launched back in April, SUSE was the first distro available for that, so it's unsurprising that they've trumped the competition again with SUSE 9.0. The Athlon 64 version is available in a separate box, which is a minor annoyance, and actually costs £25 more, which could be annoying for some users – we're not sure quite what is in the box to justify the extra cost, but, as it's the only thing out for Athlon 64 at the time of writing (although Mandrake 9.2/64 is due out soon), SUSE can pretty much charge what it pleases.

Moving on to the documentation, SUSE 9.0 has the biggest ever manual seen in SUSE Linux, which is saying a lot given that the documentation in the past has always been of the highest quality. This latest tome is a massive 1050 pages long and covers topics from the mundane to the esoteric and as wide-ranging as system installation and configuration down to chatting with your friends over *Kopete*. If anyone ever complains about the cost of 'free software', perhaps SUSE should clout them around the head with this really useful manual.

SUSE 9.0 > SUSE 8.2

The sheer weight of changes made in SUSE 9.0 make it a very viable upgrade from 8.2, particularly if you're a long-standing SUSE user. If you're using another distro and are thinking of switching, the choice will take a little more consideration – SUSE is not as easy-to-use as Mandrake, despite all the new effort that has gone into SUSE 9.0, however on the flip-side it is much more stable once you get it working. We're confident the installation problems we had won't be seen much in the outside world, because a lot of the machines we tested it out on were two or three years old – not exactly the latest kit!

It will be interesting to see how SUSE's new branding strategy pans out, as SUSE 9 was produced and finalised before the new branding was made public. As a result, references to "SuSE" (note the lowercase 'u') are throughout, alongside the old-style chameleon and a 8.2-like KDE theme.

The new software included with SUSE 9.0 makes it a competitive upgrade for some existing users, however you do need to question whether it's worth laying out cash just to get the latest versions of what are otherwise free packages. Yes, the manual is as good as you will find, and the system is fairly rock-solid once sorted out, but given the headline change in SUSE 9.0 is the inclusion of OpenOffice.org 1.1 – something SUSE 8.2 users can download and use for free – it really shows that even SUSE mustn't think a great deal has changed on its behalf.

Nevertheless, this is the best SUSE distro available, and also continues to arguably be the finest all-round Linux distribution available. While it has its problems, these are easily outweighed by the easy install and administration, comprehensive documentation, and large selection of packages. **LXF**

LINUX FORMAT VERDICT

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

A little troublesome at first, but it's well worth the hassle to get this working so you can enjoy all the latest apps and utils that make SUSE number one.

RATING 9/10



HARDWARE

Linux input devices

PART 2 Gaming devices are a lot more complicated than the Linux-supported keyboards we examined last issue, as **Paul Hudson** found out...

While there are thousands of games available that make use of devices such as joysticks, steering wheels, and joypads, for one reason or another many

people still have the perception about it that linux supports gaming badly. This is a shame, because with the rise of commercial-quality drivers such as nVIDIA's we find ourselves in the position where games *can* be run

perfectly well. Playing a game, however, needs more than just fancy graphics – if the sound doesn't work properly, or, more crucially, if your input device of choice isn't supported, then you're in big trouble!

The situation with joysticks etc is complicated because there isn't any standardised input format for these devices behind the scenes, so each device usually needs to be handled in its own way. While this might not seem a problem, what it eventually means for such devices is that the amount of popularity a product gets dictates how likely it is to get a working Linux driver, with some products not even getting a driver at all. Undaunted, we picked four of the top gaming devices currently available on the market, and gave them all a whirl using Debian.

Logitech MOMO wheel

BUYER INFO

- **SUPPLIER** Logitech
- **PRICE** £100
- **WEB** www.logitech.com

LINUX FORMAT VERDICT

FEATURES	8/10
EASE OF USE	N/A
VALUE FOR MONEY	8/10

A great piece of hardware, but with no direct Linux support available – yet!

RATING **6/10**

Logitech has long ruled the domain of input devices, and so it's no surprise that this is the first of three Logitech products here! This steering wheel is quite brightly coloured, with an attractive design modelled after the famous MOMO racing brand. Also sporting a stick shift, it also has the more conventional F1-style paddle-shift that is seen in other steering wheels. Like most modern devices this wheel runs entirely through USB, which should make it fairly easy to work with. One minor niggle is that the wheel attaches to your desk using a large plastic screw, which didn't hold the unit on nearly as firmly as we would have liked.



Just plugging the device in didn't work with all the input drivers enabled, which usually means you have to look elsewhere for support. We hunted high and low for info on working drivers and the problem seems to be fairly straightforward – the Logitech Formula Force device, a wheel very close to the MOMO, has working support so we've

been told, and if that can work then it's only a small leap from there to getting the MOMO to work. However, without a working driver, we couldn't get the thing registering its input for our games to work with, which means this device just isn't going to work for the time being.

MS Sidewinder Force Feedback wheel

BUYER INFO

- **SUPPLIER** Microsoft
- **PRICE** £100
- **WEB** www.microsoft.com



It was Microsoft who pioneered force-feedback equipment on the PC, so it's unsurprising that their wheel has long been respected for its high quality. As soon as the wheel is powered on the force effect kicks in with auto-centring, and it's quite obvious that it has a very powerful engine inside for the force

mechanism. Like the MOMO, the Sidewinder runs entirely from USB, and has paddle-shift gear changing behind the steering wheel. However, unlike the MOMO, the Sidewinder has a firm screw to attach it to your desk.

Getting the Sidewinder to work was the easiest of all the devices here, as the drivers came bundled with our OS and just needed a little configuration – it was only minutes from plugging the device in to playing games. As you may know, Linux has had a driver for this wheel back as far as kernel 2.2, mostly thanks to the work of Voltech Pavlik at SUSE. Naturally Microsoft doesn't support Sidewinder on Linux, but there are other avenues!

Playing games on the Sidewinder is great fun because it's a well-designed

wheel that fits your hand nicely. The reason for this device's Linux compatibility is down to the fact that it's a very popular peripheral; and the simple fact that installation is such a doddle, probably even for a newbie, makes this a great choice for any Linux gamer's hardware setup.

LINUX FORMAT VERDICT

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

Installed in minutes, this is a great wheel that will give many hours of fun.

RATING **9/10**

Logitech Cordless Rumble Pad

BUYER INFO

- **SUPPLIER** Logitech
- **PRICE** £40
- **WEB** www.logitech.com

While force feedback is prevalent and cordless peripherals have been popular for a while, it's been a case of "never the twain shall meet". Into this market the Cordless Rumble Pad steps, proving that a mix of the two is certainly

popular, and can still be squeezed into a fairly small unit.

For driver support, the rumble pad is experimentally supported by the HID driver, which means it's fairly easy to configure on paper. It's currently the only force-feedback gamepad to have such driver support, which certainly sets it apart in the field! Despite the rumble pad support being marked experimental, we found it worked first time with no special effort required.

LINUX FORMAT VERDICT

FEATURES 8/10

EASE OF USE 9/10

VALUE FOR MONEY 9/10

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

RATING 9/10



Logitech Freedom 2.4GHz

BUYER INFO

- **SUPPLIER** Logitech
- **PRICE** £50
- **WEB** www.logitech.com

With the popularity of wireless devices continuing to storm ahead, wireless joysticks such as this one are starting to become more popular. Operating at 2.4GHz should mean you should experience no lag when using the stick, which in turn should limit your blaming possibilities when you lose! Logitech has years of experience

making joysticks, and this shows in the design of this unit – it fits your hand well regardless of whether you are

LINUX FORMAT VERDICT

FEATURES 8/10

EASE OF USE 10/10

VALUE FOR MONEY 7/10

Works with Linux, yes – but under that styling it's a very traditional passive design.

RATING 8/10



right- or left-handed, and the buttons are all within easy reach, whatever the size of your hand. On the downside, the price is quite high – how highly do you prize the lack of wires?

On the driver front, this again seemed to work fine using our default configuration. This is partially helped by the fact that this isn't a force-feedback stick, but also because it's not substantially different to Logitech's other joysticks – the wireless aspect is completely irrelevant at the driver level.



Conclusion – Linux input devices

Getting advanced controllers such as these to work on Linux is still not an easy job by any means – even here at LXF Towers we had to expend a considerable amount of effort to get these devices working as much as we did. While it's possible – perhaps even probable – that users running an 'easy' distro such as Mandrake will have a simpler time than us; at the end of the day it's still not going to be as easy as we'd all like. What's really needed is a unified driver structure for these devices that make them 'just work', and while we're not all that far away from this currently, there are a few small holdbacks:

- i Proprietary code. Some manufacturers, including Logitech, use proprietary code by Immersion Corporation for their force-feedback support. This is not released under the GPL, which means either we have to try and figure it out ourselves or we get Logitech to change.
- ii Game support is still patchy. SDL has come on leaps and bounds recently, but still many people choose to write games based on their own custom toolkits. While this is not a bad thing *per se*, it does make the job of writing a driver that bit harder.
- iii Distros configured for home use don't help users install hardware. It's

fairly simple to detect hardware being plugged in, and also simple to look up vendor and product IDs against a list – at the very least a desktop distro could say "You've plugged in XYZ, and you now need to reconfigure your kernel to enable ABC".

iv Manufacturers show little willingness to change. If you can't program, hate writing documentation, and don't want to spend the time helping answer newbie questions on IRC, you can still help the Open Source movement by writing to Thrustmaster/ Logitech/ Saitek/etc and asking them to help with the work on a Linux driver for their products.

Of course, by the time this article makes it out this may already have started to change – and we hope so! Until then, don't give up on your devices; use Google to see what others are saying and try things out for yourself. Very often hardware faults will keep something from working, but be sure to try out every possibility first. For example, enable the 'Long timeout' option for USB when you compile your kernel – many devices, including several Saitek joysticks, require this to be enabled. This has been removed in the 2.6 kernel and may be removed as a backport to 2.4, but it's worth checking to make sure. [LXF](#)

Extreme Programming with Ant

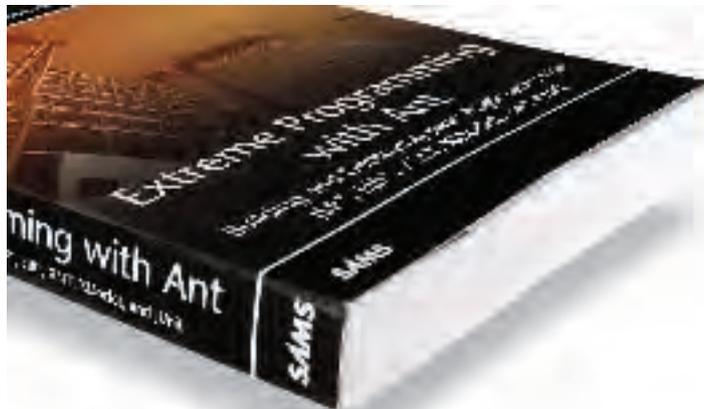
If you thought *Ant* was just a glorified version of *make*, this book might be for you, says **Paul Hudson**.

BUYER INFO

- **AUTHORS** Glenn Niemeyer and Jeremy Poteet
- **PUBLISHER** Sams
- **ISBN** 0-6723-2562-4
- **PRICE** £25.50
- **PAGES** 444

The concept of Extreme Programming (XP) has been around for quite some time now, and has managed to gain quite a few loyal followers. However, for many of us it remains just a theory that looks nice on the surface, and has yet to really make any impact on our work. *Extreme Programming with Ant* has been written with precisely that audience in mind, and, starting with the lowly *Ant*, aims to explain everything you need to for XP development on Java.

The book is ordered around the XP lifecycle, presenting the enabling Java technologies at each appropriate juncture, however the book doesn't



contain all that much theory. Of its 444 pages, we reckon that only about 10 pages are theory – not necessarily a bad thing, because there are already a great number of books that are far too theory-heavy, so it's a relative breath of fresh air to see an XP book that's so hands-on.

Perhaps one of the key advantages to this book is that it doesn't focus on any one technology, instead spreading

the content across a variety of closely knit tools, each offering its own unit advantages for XP zealots. For readers looking for a good all-round introduction to the tools that support Java programmers, this book actually – although perhaps inadvertently – fits the bill rather nicely. For example, the chapter on XDoclet presented a lot of information entirely new to us, which alone made the book a worthwhile read.

The major downside to the book is that it does cover a number of technologies that most people aren't likely to want to use. If we were to be charitable, we'd say that perhaps only about 50 per cent of the book really contains information likely to be of interest to the majority of readers, although the actual figure might be closer to 35 per cent. While the rest of the text isn't uninteresting, it is probably too esoteric – it takes a really dedicated programmer to want to customise *Ant* to quite the same level the writers want to!

We'll be printing a feature all about Extreme Programming in the next issue of *LXF* to bring you up to speed.

LINUX FORMAT VERDICT

A thorough treatise on XP with Ant, JSP, EJB, and other enabling technologies.

RATING **8/10**



A Practical Guide to Curl

Paul Hudson is famed in these parts for his ability to Curl one off...

BUYER INFO

- **AUTHOR** Kevin Hanegan
- **PUBLISHER** Charles River
- **ISBN** 1-58450-288-6
- **PRICE** \$41.95
- **PAGES** 399

Although the Curl has been around for quite some time, it doesn't seem to have really caught on yet. That's not to say Curl isn't a great idea, merely that there clearly isn't enough momentum behind the movement for it to reach a critical mass. This book, from an author more commonly associated with Perl, seeks to redress the long-term ignorance about the power of Curl by teaching the language through association with similar features in other languages.

The "practical guide" part of the title is quite accurate, as all of the text is directly hands-on – you can pretty



much take large chunks of the example code and deploy it in your own projects with little fuss. However, this does have the downside that the constant need to provide copious practical examples means that progress is slow – by the time you reach the halfway mark you're still covering basic concepts such as the if statement and the switch statement.

Given the relative shortness of this book, we were quite surprised to find quite how much information is covered

– you're taken on a rollercoaster tour of GUIs, animations, persistence, and sockets, all of which are covered in fairly good detail and interspersed regularly with illustrations and screenshots. One unusual but nevertheless interesting chapter is that on Game Development – quite an unorthodox topic in Curl, but it actually serves to give exciting code examples that you really want to try out yourself.

The biggest downside to the book is that it is practical in a very generic way

in that there are lots of short code examples to handle, for example, US zip code matching and other simple tasks. However, the book lacks any sort of in-depth development content where various tasks are designed and implemented in Curl over more than just a few pages. As a result, the code is easily dropped into your own work, yes, but there's little discussion on how to handle large-scale Curl development and group programming efforts. This isn't a fatal problem, particularly if you're fond of 'cookbook'-style books, and really is just a direct result of the lack of pages.

LINUX FORMAT VERDICT

An excellent, if somewhat short, reference guide to most of Curl. It could use another hundred pages or so for catering to the needs of more demanding readers.

RATING **8/10**

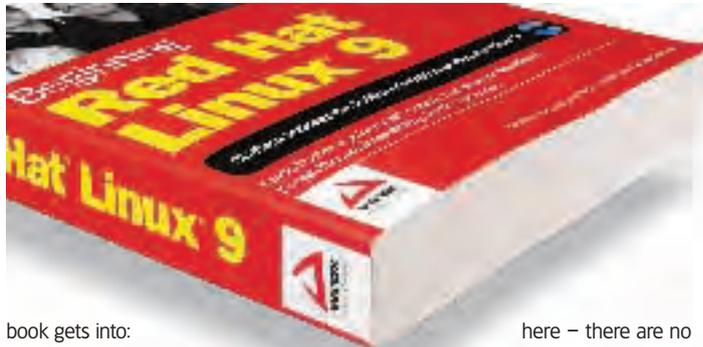


Beginning Red Hat Linux 9

Andrew Hudson comes in from the cold and discovers the delights of Linux and Red Hat 9.

BUYER INFO

- **AUTHORS** Sandip Bhattacharya, Pancrazio De Mauro, Mark Mamone, Kapil Sharma, Deepak Thomas, Simon Whiting, Shishir Gundavaram
- **PUBLISHER** Wrox
- **ISBN** 0-7645-4378-4
- **PRICE** £26.99
- **PAGES** 570



When this book dropped onto my desk with a note attached saying "Review me", I thought it would be an excellent opportunity to get back to basics and look at Red Hat from a beginners point of view. There are dozens of books on the subject and it's always interesting to see how each book introduces Linux and Red Hat.

Skimming through it, the book seems to be like any other beginners book – lots of pictures and easy-looking summaries. However, a detailed inspection shows the real detail the

book gets into: within the first six chapters you are not only introduced to the graphical side of Red Hat, but also taken through shell commands, all of which is described in often painstaking depth to make sure readers fully understand each topic. It even touches on shell scripting towards the end of chapter six.

The writing style is actually quite consistent, considering that seven people were involved with it. If you compare this style of writing to the *Dummies* series of books, you won't find the same light writing manner

here – there are no cartoons and off-topic jokes, for example. It reads more like a school textbook, which is no bad thing, because the book still manages to get the message across in an easy-to-understand way, giving you plenty of examples to try out yourself. Unfortunately, there are some niggling layout problems that should have been caught by the proofreaders, but that doesn't detract from the overall quality of the text.

As the book goes on, it gradually gets more and more in-depth, and the later chapters cover setting up

networks, managing the system with Perl, and security. The book is wrapped up with a superb chapter exploring further things to do with Linux, giving details of where to go for resources online, examples of system admin applications like webmin, GNOME workstation command center as well as a quick look at the various programming languages available for Linux. It certainly gave me a few ideas of where to go next.

If you are new to Linux and looking for a book that assumes you have just a little computer knowledge, you can't go wrong with this book. All that lets it down is the odd slip in the layout, but aside from that it's suitable for anyone who wants to get started with RH.

LINUX FORMAT VERDICT

Packed with info, this is a great introduction to Red Hat 9 and Linux in general that gently introduces advanced topics.

RATING **9/10**



Who will keep you on the Information Highway?

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

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

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

Hostway - straightforward Linux hosting.



A FREE MONTH
of web hosting.
When ordering online
use "FUTLIN047".

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

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

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

HOSTWAY
THE HOSTING COMPANY

Mastering Red Hat Linux 9

Richard Cobbett risks a hernia as he carries this hugely comprehensive guide over to his desk!

BUYER INFO

- **AUTHOR** Michael Jang
- **PUBLISHER** Sybex
- **ISBN** 0-7821-4179-X
- **PRICE** £37.99
- **PAGES** 942

Many reference books can be fairly be described as doorstoppers, but *Mastering Red Hat Linux* makes a solid play at being an actual, literal door. It's a huge tome, starting at the moment you slip your CDs into the drive and not stopping until you're fully indoctrinated in its penguin-stamped ways. As with many Red Hat-focused books, you get a full copy of the Publisher's Edition to play with immediately – this is a cut-down version of the main Red Hat distribution that won't be adequate for users to attempt running any servers on, but does give readers all of the main applications that they need to give an adequate demonstration of what Red Hat can do.



As ever, the bulk of the book focuses on using Red Hat from the GUI instead of a command line, although in a slightly more focused method: using the graphical tools for as much as possible and waiting until the final chapters to get into **cd** and the other main commands, with a full reference guide to navigation, partition managing and configuring networks with just a handful of keypresses.

This doesn't however mean that *Mastering Red Hat Linux* skips the

tougher elements. It describes itself as a starter guide to both desktop and server machines, and this is precisely what it offers – with a primer on TCP/IP settings and configuration, guides to managing Linux on your LAN and using *iptables* and *Etheral* to secure your network after you're done. This done, it moves on to more specialised packages, including bridging the Linux/Windows gap using *Samba*, configuring and troubleshooting *Apache*, *CUPS/LPD* and even...deep breath...a quick introduction to *vi*.

You won't have mastered them all by the time that you hit the final page, but you will be well on the way to doing so. And, more importantly, by the time that you get that far, you will have picked up all the information that you need to get working with both GNOME and KDE, the *OpenOffice.org* suite and also the old KOffice and GNOME Office replacements, and you'll have a system capable of running anything else that takes your fancy. Get all the way to the end, and there will be little to hold you back from achieving the title of the book – and the organisation of the book makes it fairly easy to use as a reference once you've become more adept at using Red Hat.

LINUX FORMAT VERDICT

Packed with information, and an excellent way to get started – whether or not you already have Red Hat installed.

RATING **9/10**



Essential CVS

Nick Veitch tries to be co-operative...

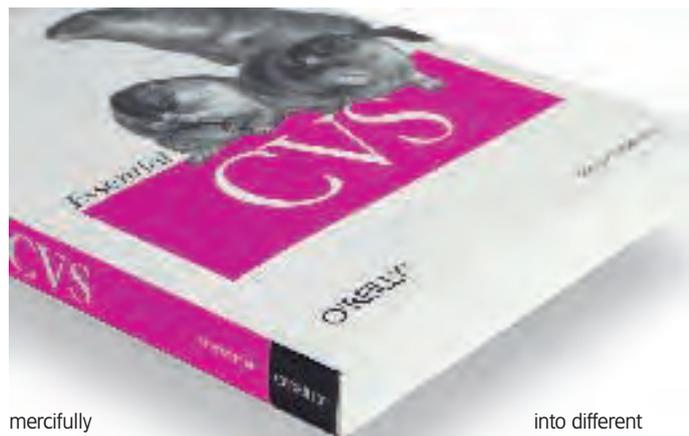
BUYER INFO

- **AUTHOR** Jennifer Vesperman
- **PUBLISHER** O'Reilly
- **ISBN** 0596-00459-0
- **PRICE** £28.50
- **PAGES** 320

The nature of Open Source programming projects generally leads to development by widely distributed teams of people, rather than individuals. In addition, these teams are often made up of people geographically separated by huge distances, or perhaps just a few doors down the road. In any case, group development of code, even on a fairly simple level, demands some sort of group management utility.

Fortunately for all concerned, CVS does just this and a lot more besides.

An introductory chapter covers the usual questions of 'What is CVS?' and 'Why do I want to use it?' This is



mercifully brief but useful for those who need convincing. Subsequent chapters than cover the basics of CVS before moving on to more advanced topics such as administering CVS servers, repository mirroring and configuring different access methods.

There's a whole chapter on tagging and branching, which usefully goes

into different aspects and strategies for managing your project too, rather than sticking simply to the mechanics of what can be achieved. A highly useful set of appendices even cover development tools and other software that can make use of CVS.

As you might hope, the writing throughout the book is clear and

precise, and there are plenty of hints and tips sprinkled liberally throughout.

It's hard to imagine some question you may have about CVS or its use which is not answered somewhere in this book, and even if you are fairly expert on CVS, there's bound to be some extra tips and tricks you can pick up here. If you just need a quick reference to CVS, O'Reilly also publish a substantially cheaper pocket guide to the subject, which is well worth considering.

For more information on CVS, read this month's *What on Earth...* feature on page 60. [LXF](#)

LINUX FORMAT VERDICT

A pretty comprehensive and useful guide to an often under-utilised group development tool that defines the way many Open Source projects are managed.

RATING **8/10**



Roundup

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

OUR SELECTION AT A GLANCE

- Slrn
- Pan
- Thunderbird
- KNode
- Gnus
- Sylpheed
- Tin
- Knews
- MyNewsGroups:
- Mahogany

Newsreaders

With his asbestos suit at the ready, **Mike Saunders** dives into the world of Usenet and newsgroups with a look at Linux's best newsreaders...



Pre-dating the World Wide Web's invention (and the Internet's commercialisation as a whole), Usenet is the grandfather of online discussion systems: mature, wise and with a few memorable anecdotes up its sleeve. Many historical gems such as Stallman announcing GNU, Torvalds requesting help for his hobby kernel project and Volkerding's first *Slackware* message can be found in the Usenet archives, along with posts on every other subject under the sun. And despite more recent popularity of mailing lists and web-based message forums, Usenet remains as strong as ever.

It's not all flamewars, though. In the right places, Usenet is still one of the best ways to engage in meaningful and enlightening discussion. Equally, as a support method it's a reliable option –

groups exist to help readers with their programming, Linux configuring, beer-making and hamster breeding. Just post a query and wait for a response.

If you're unfamiliar with newsgroups, it may seem like archaic technobabble at first glance. There's certainly a great deal of history and culture behind this; just as there are unwritten rules and etiquettes to which one should adhere on mailing lists, newsgroups also have their own standards of behaviour. Flamebaiting is out. Asking a question without doing a little research beforehand is a bad idea. *Ad hominem* insults are also definitely a no-no.

Essentially, Usenet is a hierarchy of 'newsgroups' (specific to certain topics) in which readers post messages and read responses, similar to email lists. However, Usenet news

servers are spread all over the globe and provide 'feeds' of posts to one another – it's decentralised, and users of one server (or 'site') may not see every group. By far the most common provider of Usenet feeds are Internet Service Providers; your ISP's support pages should have information on their news server's URL.

Once connected, you can browse through the list of groups and choose ones that meet your interests. With 30,000+ groups (and always growing) though, finding the most relevant can be a bit of a nightmare, but <http://groups.google.com> is supremely helpful here. Usenet isn't the only network – many companies run internal news servers – but it's by far the most used. You'll need some software to effectively read, manage

and reply to groups, and in this roundup we're putting ten of the best and most popular under the spotlight. Various notable features worth looking out for include:

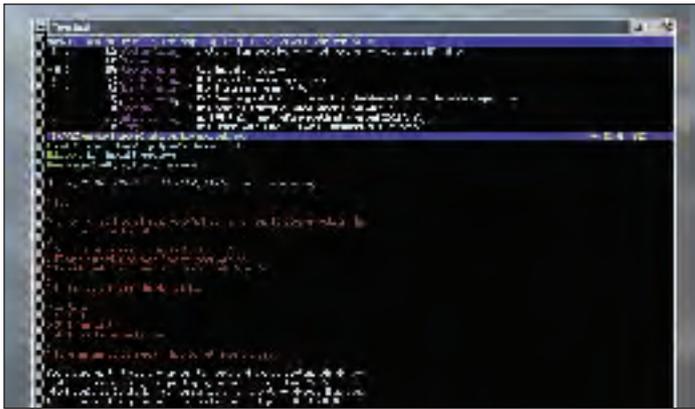
- Filtering to avoid unimportant/offensive posts
- Offline mode (download, disconnect and read on dialup modems)
- Image viewing and binary downloading for filesharing groups
- Decent keybindings (shortcuts) for power users
- Spelling checker to avoid making embarrassing mistakes

We last looked at newsreaders back in *LXF07*, so if you're unfamiliar with newsgroups, a recent Linux convert or just fancy a change from your usual app, this *Roundup* will help you find the best for your needs.

Slrn

One of the old stalwarts, with a hardcore fanbase.

■ **VERSION** 0.7.2 ■ **WEB** <http://slrn.sourceforge.net>



A split two-pane system is used in *Slrn*'s main navigation view.

As Usenet has been in action for over two decades, it's not surprising that some clients have a lengthy history. The full details of *Slrn*'s past have been lost

in time; though the earliest post from this client can be dated back to 1994, how long it had been in development before then is still anyone's guess. *Slrn* is

still in wide use by long-time Usenetters and has a reputation for being rock-solid and flexible.

Most recent distros provide a binary package of *Slrn* in their archives – if not, compiling from source is a breeze as it's purely text-based and doesn't rely on any graphical toolkits. You'll need *Slang* (*slang-devel* in many distros), and the usual **.configure**, **make** and **make install** routine will be enough. Versions also exist for BeOS, OS/2, Windows and even the Amiga.

Getting *Slrn* running requires a few hoop-jumps – the news server's location needs to be set with the NNTPSERVER environment variable, and then a *newsrc* file created with **slrn -f ~/.jnewsrc --create**. Finally, you'll need to make an `~/slrnrc` by copying over the default. *Slrn*'s colourful interface is improved by the *Pine*-esque key list along the bottom, and after an hour or so of use, navigation becomes second-nature.

By default, *Slrn* uses *vi* for message composition (spellchecking delegated to the editor), but this can be changed to alternatives such as *Emacs* or *Pico*. If the binary hasn't been linked with the

uudeview library, its internal routines suffice for downloading binaries but naturally there's no inline viewing of images. Coupled with the *slrnpull* utility, *Slrn* offers a workable offline reading mode (along with message postponing), while the Scoring system excels for creating immensely versatile filters.

With its text-based interface, extensive key command set and steep learning curve, *Slrn* isn't ideal for everyone. Conversely, experienced users will love the flexibility it provides, and the in-depth docs are first-class. It's certainly faster than its graphical siblings – long, slow development cycles have resulted in sheer stability.

LINUX FORMAT VERDICT

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

Undeniably thorough in its options and featureset, and rock solid too.

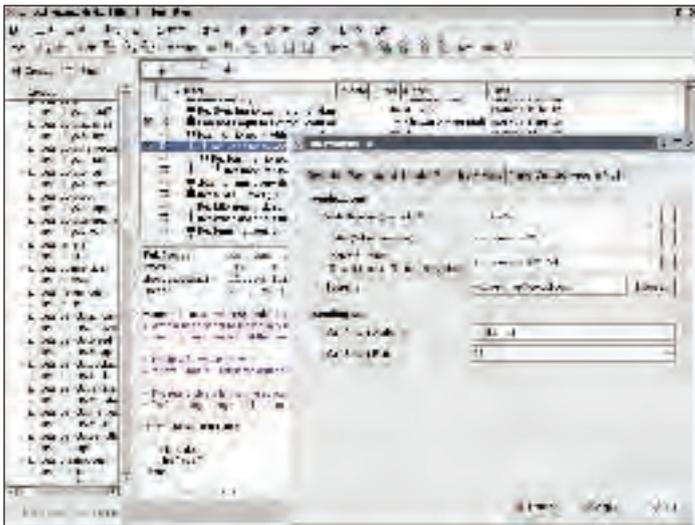
RATING **8/10**



Pan

A lot more than just a flash in the pan...

■ **VERSION** 0.14.12 ■ **WEB** <http://pan.rebelbase.com>



Pan's attractive layout and icon set, with the Prefs box in front.

Like *Enlightenment*, *Gaim* and *Firebird*, which are all very solid for day-to-day use, *Pan* is one of those apps with a low-ish version number

that doesn't reflect its stability or featureset. *Pan* has been in development for several years, picking up many fans along the way, and has

been covered in *LXF* a few times before – most recently in issue 42's *Hot Picks*. Various features and bugfixes have been made in the meantime, and here we're looking at version 0.14.2.

You'll need the appropriate development libraries and headers for GTK2 (2.0.5 and above), along with *libxml* and *gnet*, if you're compiling *Pan* from source. Fortunately, these are the only major dependencies and should be installed (or at least readily available) in most distros. Optionally, *Pan* can also use *gtkspell* for the spellchecker functionality. We've provided RPMs for Mandrake and Red Hat on the coverdisc along with the source.

When first started, *Pan* pops up a slick and friendly wizard to configure the program's settings; this requests the news and mail server details together with personal info. Using a resizable three-pane layout, clean toolbar icons and informative status bar, *Pan*'s interface is accessible to newcomers (although the plethora of menu options could be a tad daunting for some users).

Among its excellent array of features, *Pan* sports multiple-server connections, inline image-viewing for binary attachments (*uu*, *base64*, and

yEnc formats), an outrageously comprehensive filtering and rule system (alerts you on matched posts) and heaps more. Article caching and offline reading and postponing is competent too. Just about every aspect of the client can be fine-tuned – external mail/browser apps, message colours, fonts and layout – but while a thorough set of keybindings is available for power users, there's no way to remap them in the app.

Overall, recent *Pan* releases could warrant a full section in this magazine to cover all the features and options in detail. It's a firm favourite of hardcore news users, yet at the same time it provides a friendly front-end for first-timers and is both speedy and robust to boot. It's on the coverdisc along with all the others here, so give it a try!

LINUX FORMAT VERDICT

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

Friendly, fast and flexible, *Pan* is the bee's knees of news clients.

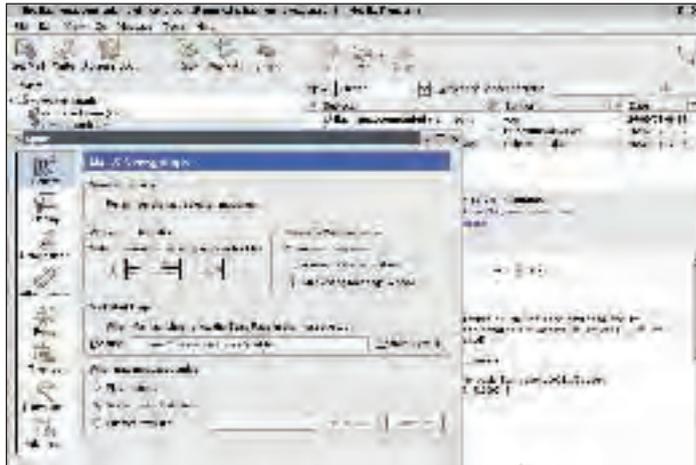
RATING **9/10**



Thunderbird

Mozilla-based mail/news companion to Firebird.

■ VERSION 0.3RC2 ■ WEB www.mozilla.org/projects/thunderbird/



Thunderbird's sweet *Qute* theme is the default. Options box popped up.

It took a long time coming, but Mozilla's 1.0 release was celebrated by pretty much everyone in the free software community – finally, a jam-packed standards compliant browser

to replace the ugly and volatile Netscape 4.x suite. The only problem was bloat; many agreed that it would be best to split the WWW and mail/news components into separate

apps, and *Firebird* and *Thunderbird* appeared. We looked at the latter's mail handling ability in LXF41's *HotPicks*, and with the update to version 0.3RC2 we're concentrating on the newsreader side here.

Building *Thunderbird* from source is not for the faint-hearted, but the project's supplied binary packages are thankfully well-produced. It'll happily co-exist with *Mozilla* and *Firebird* installations, dropping its config files in `~/.thunderbird`, so no problems there.

Like *Pan*, *Thunderbird* throws up a comfortable wizard on first start which steps through basic configuration and account setup. The crisp *Qute* icon theme works well – others are also available – and particularly sweet is the on-the-fly group searching. Just type in a few letters of a group name/topic and *Thunderbird* will show the matches. Top notch.

A spelling checker is available, while multiple server support, inline image viewing and a simplistic filtering system are also present. Keyboard shortcuts have been added for the most common operations, but they won't satisfy power users. *Thunderbird's* design evidently targets

the middle-ground between newcomers and experienced users, and of all the clients on test here it's the easiest to get started with.

Ultimately, *Thunderbird* suffers in one crucial area: performance. It's slow. Not unusably so, and on modern machines it runs at an acceptable lick, but compared to *Slrn* and *Pan* it feels cludgy and rough. These are early days and if the client sees the same improvements its sibling *Firebird* did, things will get considerably better. Speed issues aside, it's pleasant-to-use with a healthy range of features and expandability through Extensions, and if you've got a fast machine it's an ideal companion to *Firebird*.

LINUX FORMAT VERDICT

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

Very pleasant for newcomers, but performance is a large problem in current releases.

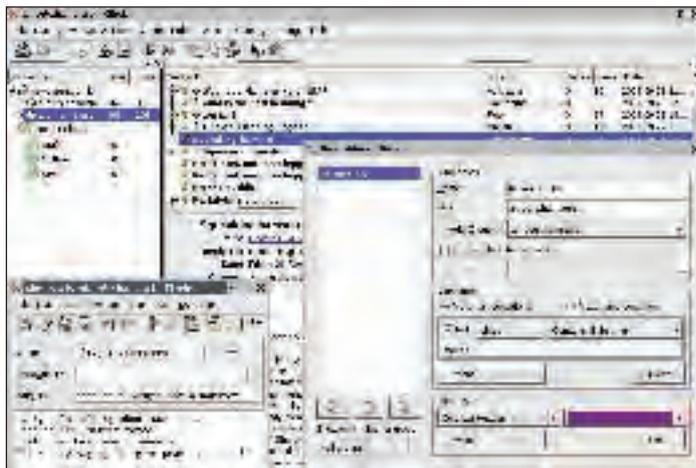
RATING **7/10**



KNode

The official KDE news client.

■ VERSION 0.7.2 ■ WEB <http://knode.sourceforge.net>



KNode's appropriately KDEish front-end, with a composition window open.

Right from the project's early days, KDE's developers intended to make the desktop environment fully Internet-equipped with a suite of tools for

emailing, browsing, chatting and reading newsgroups. And this they've done marvellously – *KMail* is highly regarded as one of the best mail clients around,

Konqueror is ever-improving with help from Apple, and *KNode* has a substantial following too. Long-time LXFers may remember *KRN* being KDE 1's official newsreader, but when KDE 2 appeared it was replaced by *KNode*.

Currently, *KNode* is supplied as part of the *kdenetwork* package and as a result it requires KDE libraries and supporting files to run. The latest release (0.7.2) is supplied with KDE 3.1, so your distro is likely to include it. If you're running a different desktop, you won't need the whole works – just *Qt*, *kdelibs* and *kdenetwork* should be enough.

KNode's GUI follows the established three-pane convention (resizable), and when first fired up it uses the Prefs box to acquire server information. As with *Thunderbird*, the search as-you-type system for subscribing is a useful touch, and overall the interface is polished and approachable and fits in perfectly with the rest of KDE.

Multiple servers are supported, as is inline image viewing, and posts can be checked for spelling errors before sending. Splendidly, the keyboard shortcuts can be redefined in the program itself, while a detailed scoring and filtering system (regex matching

supported) has been implemented too. The only major feature that's missing is an offline reading mode; this drastically reduces usability on dialup connections, and while the coders suggest using the *Leafnode* tool to rectify this problem, it'd be much better to have the functionality in the program itself.

Nonetheless, *KNode's* performance is respectable, the documentation is first-rate and we encountered no stability problems at all in several days of heavy use. It clearly works best in its native desktop environment, but if you're not a *Pan* fan and want an alternative graphical client with a clean design and sizable range of common features, it's a wise choice.

LINUX FORMAT VERDICT

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

The closest to *Pan* in terms of overall polish and features, and another treasure in KDE's Internet suite.

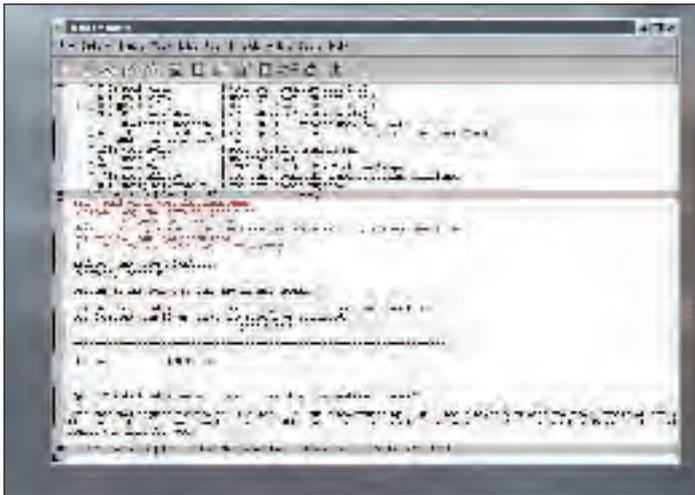
RATING **8/10**



Gnus

There's a *Gnus* loose about this hoose!

■ **VERSION** 5.10.2 ■ **WEB** www.gnus.org



Emacs in X11-mode, with its cram-packed two-paneled browsing view.

Emacs: an 'extensible, customizable, your view (and usually whether or not self-documenting real-time display editor' or 'Eight Megabytes And Constantly Swapping', depending on you're a *vi* fan). Arguably the most famous text editor in existence, *Emacs* has an army of dedicated supporters

and is one of the GNU project's many successful applications. Its LISP engine is indubitably powerful in expanding the editor for other purposes, and one such extra is *Gnus* (or in recursive fashion, 'Gnus Network User Services') a LISP mail and news client.

Many distributions, including Mandrake and Red Hat, include *Gnus* as part of their core Emacs packages. If not, you should be able to find it in an external package (eg *emacs-gnus* or similar) or you can download it from the project's main website. As it's normally distributed in bytecode-compiled LISP, there should be no binary compatibility problems and the only dependency is a working *Emacs* installation.

Starting *Gnus* successfully requires a bit of fiddling – the easiest way is to set your news server in the NNTPSERVER environment variable *a la Slrn*, launch *Emacs* and run **M-x gnus**. The client will choose a few groups by default; to begin reading, enter **gnus-fetch-group** followed by the group's name. There's a lot more to it than this, but the info pages (*info gnus*) hold a proper tutorial.

As with the other text-based apps in this roundup, *Gnus* races along and its extensive development history

(going back to 1988) has led to a very reliable and flexible client. The *Emacs* spelling checker can be used in message composition, the default keybindings are love or hate (Ctrl used to be where Caps Lock is now on older keyboards) and multiple server support is available. Offline reading can be done with some work – see the *Gnus Agent* documentation – and scoring/filtering is equally in-depth.

All things considered, *Gnus* is only really an everyday option if you're already an *Emacs* user; it's tricky to work with initially, and *Slrn* is a better choice if you just need a text-based client. Despite the somewhat threatening interface though, it's remarkably well documented.

LINUX FORMAT VERDICT

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

A recommended choice for experienced *Emacs* fans, but otherwise too steep learning curve for casual users.

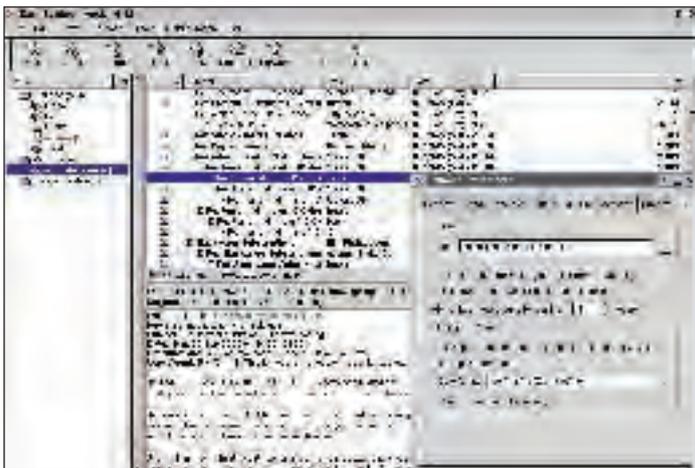
RATING **7/10**



Sylpheed

Fast and slimline combined mail and news app.

■ **VERSION** 0.9.6 ■ **WEB** <http://sylpheed.good-day.net>



Sylpheed's GUI isn't quite as slick as its GTK2-based rivals.

Mail managing and newsreading tasks share enough similarities that they can be effectively combined together in a single app. *Thunderbird* does this competently, as does *Pine* (though the latter's news handling code is somewhat lacking), and *Sylpheed* is heading in the same direction. At the

moment, development is split into two branches (the stable one on test here and *Sylpheed-Claws* for bleeding-edge goodies) and although the lead developer is Japanese, there's plenty of well-translated documentation.

Sylpheed is built around the GTK+ 1 toolkit; this may be fairly old now but all distros still include it to support earlier apps, so it should compile without hassle. Having *imlib* installed is also a good idea – it provides inline image viewing – and the RPMs require GnuPG and GPGME so we've included those on the coverdisc too.

Because it doesn't yet use GTK2, *Sylpheed's* interface doesn't sport fancy anti-aliasing and other frills, but it's clean and uncomplicated with a fitting icon set. The Account Prefs box handles all mail and news server settings, allowing multiple servers, and our only gripe with the front-end is that it doesn't provide much feedback during operations.

Like *Thunderbird*, *Sylpheed* doesn't offer any single-key shortcuts (instead relying on CTRL+ modifiers) but most functionality is accessible through the keyboard. On the feature front, inline image viewing (there are some

problems with unencoded files) is supported, but there's no spellchecker unless an external editor is used, nor is there any set offline-reading capability.

Even with these problems taken into account, it's apparent that *Sylpheed's* main focus is on the mail handling side and subsequently the NNTP code and options shouldn't be judged too harshly. Feature-wise it's lacking compared to *Pan* and *KNode* – power users will find it too limiting and newcomers will prefer the friendly approach of *Mozilla*. But if you want to combine your mail and news clients and deal with it all from a single app, it's unquestionably much lighter and speedier than *Thunderbird*.

LINUX FORMAT VERDICT

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

Not brilliant as a standalone newsreader, but works well if you want a mail client tied-in too.

RATING **7/10**



Tin

Text-based classic which doesn't munch resources.

■ **VERSION** 1.6.2 ■ **WEB** www.tin.org



Colours help readability in Tin's Elm-esque design.

Tin rivals Gnu and Slrn in its age and development history, with the first release appearing way back in 1991. Consequently, it's not as popular today but many Usenet old-timers still rate it highly; one reason in particular is its simple text-based interface, which bears a striking resemblance to the equally venerable Elm mailer.

To get Tin started quickly you'll need to pass a few flags on the command line:

```
tin -r -g <news_server>
```

should do the trick, but you'll also

need to delve into the configuration files stored in ~/.tin for best results. Although the list of keybindings along the bottom helps, it's the hardest of all the clients to get to grips with (try the -a flag for a more colourful display like the one in the screenshot).

Tin's extraordinary range of options and configuration settings makes it hugely versatile, although the lack of a built-in offline reading mode and hard barrier to entry don't help. Be prepared to read the man page thoroughly first!

LINUX FORMAT VERDICT

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

A bit too taxing for normal newsreading jobs, but very light and robust.

RATING 5/10



MyNewsGroups:)

Read and search newsgroups in your web browser.

■ **VERSION** 0.6b ■ **WEB** <http://mynewsgroups:.sourceforge.net>

Web-based email systems have been taken up by companies as they provide a consistent front-end regardless of OS and client, and improved manageability. MyNewsGroups:) (the smiley is officially part of its name) aims for a similar goal, but for NNTP instead of mail.

It's built in PHP4 on top of a webserver, with the MySQL database used as a store, and provides an attractive and workable front-end (admin'd through the browser too). Users can search through messages and check out stats, and already some large sites have put it into place as an easy-going help resource.

Naturally, MyNewsGroups:) doesn't compete directly with the other newsreading tools in this Roundup and

individual users won't be interested, but if you run a popular site and want to provide your readers with a sane and simple method for browsing Usenet, it's a superb solution. Best of all, the smiley turns into a sad face when an error occurs!

LINUX FORMAT VERDICT

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

Not meant for standalone use, but definitely one of the best web-to-NNTP gateways around.

RATING 8/10



Knews

Bright yellow widget sets? Lovely...

■ **VERSION** 1.0b.1 ■ **WEB** www.matematik.su.se/~kjj/

Whereas Thunderbird, Pan and KNode all require various libs, toolkits and other dependencies to run, Knews opts for a vanilla X GUI making it more suitable for older machines. It's easy to build (providing you follow the xmkmf path as described in INSTALL) and starts up just as quickly as the text-based clients.

When you launch Knews, it starts up to a blindingly garish creamy-yellow interface. Aesthetics aside, Knews is a breeze to operate – just tap in the news server and groups you prefer and then read away. There are keybindings for the main operations, but no inline images, multiple server support or inbuilt spellchecking functionality, which is a shame.

It doesn't compare to the other graphical clients in terms of features and eye-candy, but if you're running an old box and don't need KDE or GTK eating up disk space and CPU cycles, Knews is no-nonsense, fairly obvious to use and does the job with pace.

LINUX FORMAT VERDICT

FEATURES	5/10
EASE OF USE	7/10
DOCUMENTATION	5/10
PERFORMANCE	8/10

Worth investigating if you need a plain X newsreader for a low-spec machine.

RATING 5/10



Mahogany

Integrated mail and news handling.

■ **VERSION** 0.65 ■ **WEB** <http://mahogany.sourceforge.net>



Mahogany's GTK1 wxWindows layout with the options box up front.

Mahogany is a well-known and popular mail/news client, being cross-platform thanks to its use of the wxWindows UI library. This supplies a 'native' look across different operating systems, with GTK chosen for Linux apps; pleasingly, the RPM supplied on our coverdisc has it all rolled-in so dependencies shouldn't be an issue. If you're building from source, wxGTK should provide the necessary development bits.

As a mail client, Mahogany performs commendably well and integrates the newsreading aspects

appropriately. Unfortunately, this is also its downfall – the NNTP server 'folder' entry corresponds to only one specific newsgroup, so if you subscribe to many it's very messy to use.

Otherwise, it's stable and pleasant to work with, and the supplied docs are brilliantly written.

While its news handling aspects aren't a threat to KNode and Pan, it's an adequate alternative to Sylpheed if you need integrated mail and news handling. The main focus is on the mail side, though.

LINUX FORMAT VERDICT

FEATURES	5/10
EASE OF USE	9/10
DOCUMENTATION	9/10
PERFORMANCE	7/10

Multi-platform and easy to get started with, but lacking in dedicated newsreading functionality.

RATING 6/10



NEWSREADERS THE VERDICT

Linux grew up on the Internet, with coders all over the world collaborating on its kernel and supporting utilities, and much development work achieved via newsgroup posts and discussions. At the same time, the technical aspects of Usenet and NNTP in general haven't changed a great deal in the last decade or so, and it's for these reasons we still see the older clients (*Slrn*, *Gnus* etc) going strong. They may appear antiquated and unhelpful to some users used to the GUI environment at first, but like text editors and mail clients they have fans who'll always support them, and who'll sometimes get involved in flamewars.

Developing an attachment to a program is natural though, especially if one spends a lot of time in it. Few people fall in love with *Is* or hex editors, but applications in which important data is stored become more than just a string of bytes, and if a program is used for hours every day its behaviour, control and quirks soak in. And that's the key to fast and effective newsreading – being able to navigate and control a client quickly and easily. Every application on test here has the basic functionality – authentication, searching, MIME, threading etc – and they're all usable for simple tasks, but it's the polish and featureset that separates them.



Trouble installing, or want to mention an app not covered here? Visit the forums at www.linuxformat.co.uk

Bearing this in mind, *Pan* comes out top in this month's *Roundup*. Its most applaudable success has been in providing a friendly and approachable front-end for newcomers (particularly with the startup wizard), while at the same time retaining the necessary options, keybindings and features that long-time Usenetters demand. And the smooth and solid icing on the cake is its performance and reliability; we've been running various *Pan* releases on-

and-off for years, and have come across very few glitches or slowdowns. It's the only newsreader here that received a perfect score on the *Good Net-Keeping Seal of Approval* tests!

Impressive goodies

If *Pan* doesn't appeal, *KNode* is certainly on its coat-tails with impressive goodies and equally satisfying performance. It fits into KDE much better too. *Thunderbird*,

meanwhile, is worth a look if you're a big *Firebird* fan but sorely needs some work on the speed issues before it can compare to *Pan* and co. It's a shame, because the *Firebird* client has benefits of being cross-platform, supporting Extensions and offering a slick interface. Of the text-based progs, *Slrn* matches a top featureset and immense configurability with a clean and quick-to-learn front-end – and an impressive history.

TABLE OF FEATURES

We tested the newsreaders on two Pentium III boxes running Red Hat 9 and Slackware 9.1. Note: memory usage varies from system to system, and is affected by the program's configuration and other factors. Pkg size is for source.

Name	License	Interface	Pkg size	Avg mem use	Inline img	Offline	Mult serv	Spell
Gnus	GPL	Text	2.2M	10M	No	Partial	Yes	EMACS
Knews	GPL	X11	380k	8M	No	No	No	No
KNode	GPL	Qt 3	N/A	15M	Yes	No	Yes	Yes
Mahogany	Artistic	wxWin	3.8M	Partial	Partial	No	No	Yes
MyNewsGroups:~)	GPL	Web	360k	N/A	Yes	N/A	N/A	No
Pan	GPL	GTK2	2.6M	25M	Yes	Yes	Yes	Yes
Slrn	GPL	Text	900k	8M	No	Yes	Yes	External
Sylpheed	GPL	GTK1	3.1M	5M	Partial	No	Yes	No
Thunderbird	MPL	GTK2	9.8M	30M	Yes	Partial	Yes	Yes
Tin	BSD	Text	1.5M	6M	No	No	No	External

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 also provide web links if you want to make sure you have the very latest version.

If you have any suggestions for things that we should cover, email us at linuxformat@futurenet.co.uk or Mike Saunders himself at mike@aster.fsnet.co.uk

HOT PICKS AT A GLANCE

Epiphany	47
Gnumeric	48
LSH	48
Apollon	49
Peacock	49
Yawk	50
Flwriter	50
BomberClone	51
Briquolo	51

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!



WEB BROWSER

Epiphany

■ VERSION 1.0 ■ WEB <http://epiphany.mozdev.org>



Epiphany and Prefs box – note the close buttons on every tab (this is a feature that is often requested for *Mozilla Firefox*).

Epiphany, noun: 'a moment of sudden and great revelation' according to the *OED*. Naming of browsers has been a sensitive issue of late, following the battle between *Firebird* database developers and *Firebird* browser coders, and *Epiphany* has also seen some conflict – there's an Open Source game of the same name. With thousands of projects in varying stages of development on SourceForge *etc*, clashes are inevitable; still, it can make life more difficult for packagers. Whatever happens, the *Epiphany* browser is deservedly receiving much attention and may become a crucial part of Linux's desktop future...

In *LXF43's* *HotPicks* we looked at *Firebird*, the slimline browser built around *Mozilla's* *Gecko* rendering engine. Several other projects exist which utilise *Gecko* (*Galeon*, *Camino*, *Skipstone* and others) but replace the weighty XUL front-end with lighter or native equivalents. *Galeon* in particular provided GNOME users with a more

appropriate interface, but development and political troubles have led to *Epiphany* moving in as the GNOME browser of choice.

Getting it up and running might be complex depending on the version of GNOME you have installed – it's built on the 2.4 release, and upgrading from 2.0 or 2.2 is a huge job if you're compiling by hand. Check to see if your distro has released GNOME 2.4 package updates, or failing that, wait for the next round of distro releases which should include the new GNOME and *Epiphany*. *Mozilla* 1.4 is also required; earlier versions may not work correctly (if at all).

Looking good

Epiphany's interface follows established window furniture conventions for browsers: menu, navigation, bookmarks and address-entry bars up top and status line sitting at the bottom. Immediately striking is the attention put in to making *Epiphany* an integral part of GNOME – the HIG (Human

Interface Guidelines) have been followed closely, while downloaded files are opened with relevant GNOME apps. Coupled with the consistent icons and GTK2 theme, the browser's slick front-end fits in much more tightly with GNOME than *Firebird* or vanilla *Mozilla*. Toolbar buttons can be re-arranged, but sadly there's no straightforward way to enable an icons-only mode for saving screen space. Although it's off by default, the infinitely useful tabbed-browsing option is available, but again there's no option to hide the tab bar when only one tab is open, eating up more room. These niggles aside, *Epiphany's* interface is clean and responsive, with a fullscreen mode and uncluttered menu structure, and if the sidebar is reintroduced, it'll be spot-on.

As can be expected, *Gecko* performs admirably with zippy rendering, excellent standards compliance and decent handling of poorly-designed pages. It's not as fast as *GtkHTML* but copes with a broader range of sites and technologies; only the most utterly broken *IE*-specific pages will choke it. Startup time is a slight improvement on *Firebird's*, and even though the requirement of a full *Mozilla* 1.4 installation guzzles up disk space, the planned separation of *Gecko* into a standalone library should help in the future.

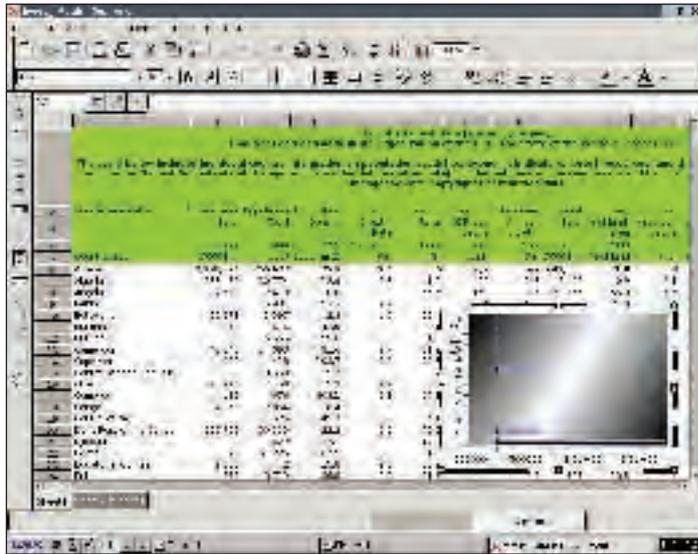
The only major issue with *Epiphany* at present, which could be a flaw or a positive step depending on your view, is the lack of options. There's no image blocking, for example, nor is there the fine-grained control over JavaScript found in *Moz* and *Firebird* – the developers are aiming for simplicity. A few extras can be tuned through *gconf-editor* and *about:config*, but as the new default browser for GNOME it needs to cater for a wider range of users.

All things considered, *Epiphany* works commendably well as a reliable and easy-going GNOME browser – it's not as featureful as *Firebird* but it certainly sports much tighter integration with the desktop. In the pipeline for the next release are support for Extensions (*ie* Flash-blockers, gestures *etc*) and refinements to the bookmarks functionality, so it's worth watching.

SPREADSHEET

Gnumeric

■ VERSION 1.2.0 ■ WEB www.gnome.org/projects/gnumeric/



Gnumeric's smooth interface with a chart being generated bottom-right.

AbiWord 2.0's arrival (see *LXF45's HotPicks*) along with the new *Gnumeric* marks the first proper release of *GNOME-Office*, which combines these two applications and other extras to make a fully-fledged office suite. *OpenOffice.org* remains the most popular all-rounder at present, but *GNOME-Office* and *KOffice* are proving to be worthy alternatives. The *Gnumeric* spreadsheet has always been one of GNOME's flagship apps, and finally version 1.2 is here.

Pleasingly, this latest release isn't tied to a specific version of GNOME and should compile and run on 2.0, 2.2 and the new 2.4 desktops (the project's site details specific library requirements). We've included an RPM on the coverdisc which should install correctly on most recent RPM-based distros; if not, building from source is generally a hassle-free process. On first start, no plugins are enabled, reducing functionality – they can be turned on in the Tools menu.

Gnumeric's most obvious visual change from 1.0 is the switch to GTK2 and a softer icon set. Otherwise, much

of the layout is virtually identical – the toolbars and menus very much resemble *Microsoft Excel*, making it easy for Windows users to navigate and operate. Similarly, the startup speed rivals *Excel* on Win32 and is still far ahead of *OpenOffice.org* (even after its 1.1 improvements).

Boasting support for all of *Excel's* worksheet functions and an additional 60 on top, along with very capable import filters for *Excel* (versions 2 up to 2000 and XP), *Lotus 1-2-3*, *Applix*, *OpenOffice.org* and the usual CSV and DIF, *Gnumeric's* flexibility is supremely impressive. Indeed, we've seen many Windows users request a native port because of its featureset and power, and that's a true sign that Linux is maturing on the desktop.

Other major updates since 1.0 include a new chart engine, in-depth text file data merging tool, performance tweaks and stacks of other touches. If you need stability, good XLS file handling, are frustrated by *OpenOffice.org's* speed or want to show your boss that Linux is ready for business desktops, try *Gnumeric*.

SECURE SHELL SERVER/CLIENT

LSH

■ VERSION 1.4.3 ■ WEB www.lysator.liu.se/~nisse/lsh/

Remote administration with Telnet is pretty dangerous over the big, bad Internet. Sending passwords in plain text is fine when you're just logging in to the headless box in the attic, or sorting out a company server in another room, but it's fatal where anyone malicious can peek and track what you're doing. *SSH* (in particular *OpenSSH*) is now the *de facto* tool for encrypted remote administration over insecure networks; after the recent security hole, though, many have been recommending *LSH*.

LSH is a client and server toolset for the *SSH* protocol (version 2). Unlike *OpenSSH* (which is BSD-licensed due to its OpenBSD origins) *LSH* is released under the GPL. Intriguingly, *LSH* has been in development for longer than *OpenSSH* but until recently the authors claimed that it wasn't ready for production use. Installation should go without any hitches – the only requirement is a working *liboop*

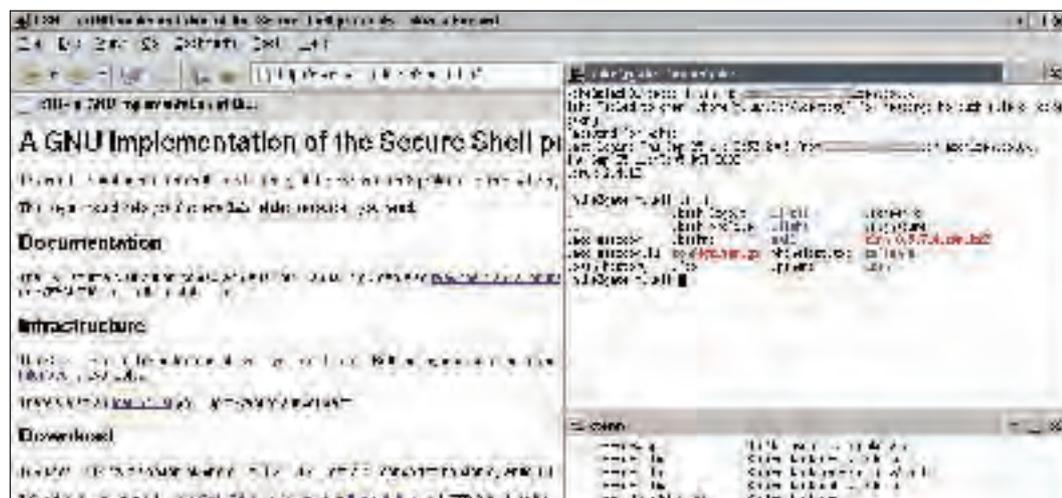
installation – on the coverdisc alongside the main source tarball.

Setting up the server involves two preparation steps: creating a seed file for random number generation and a hostkey file. From there, the server is

run just with **lshd** or **lshd --daemonistic**; there's no configuration file and all options are passed on the command line. It'd be nice to see a commented config file in future releases, as the man page is bare and outdated. AES256 is the server's default crypto method, with *Blowfish*, *3DES*, *Serpent* and others available, while public-key user authentication is also supported (and SRP uniquely).

On the client side, *lsh* by default writes keys to `~/.lsh/captured_keys` from where they can be identified and

copied over to `known_hosts` (this can be altered for *OpenSSH*-like behaviour). Features including port forwarding, compression and X tunnelling are all catered for. The *LSH* client had no problems talking to *OpenSSH* servers in our testing, and it works fine as a drop-in replacement. Similarly, the server is robust and the overall issue is security: a remote root exploit was found in *LSH* recently, so there's no guarantee that it's a better choice than *OpenSSH*, but it does deserve a look.



LSH's site, with an active session and help listing on the right.

KDE P2P CLIENT

Apollon

■ VERSION 0.8.7 ■ WEB <http://apollon.sourceforge.net>

In recent years, peer-to-peer networks have exploded in Internet popularity. As opposed to traditional systems in which many users download information from a single server, P2P networks are largely decentralised and allow anyone to share files and connect directly to others. It's a free-for-all, and some companies haven't been too chuffed; Napster, one of the first major P2P networks and a haven for music sharing, suffered as a result. Still, with millions of users and terabytes of stuff to grab, it's worth having a good client around and Apollon is one of the best newcomers.

As it's built around KDE 3.1, you'll need the relevant development libraries and headers to install from source. RPMs for Red Hat and Mandrake can be found on our coverdisc – you'll still need some other extras first. The *gift* daemon

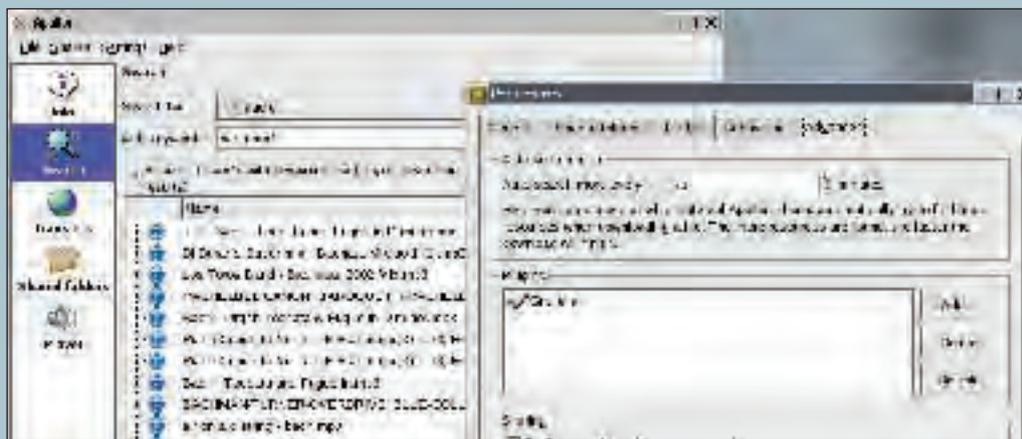
acts as a middle-man between the graphical program and a network, and supports plugins for Gnutella, FastTrack (Kazaa) and OpenNap (Napster) among others. You'll need to build *gift*, then the plugins you need, run *gift-setup*, and finally *giftd*

for the back-end. If the *Apollon* RPMs then complain about no *gift* package installed, use the `--nodeps` switch.

Despite the convoluted installation process, *Apollon* itself runs fine with a smart and polished front-end. Using a two-pane view, the program shares (no pun intended) visual layout with other popular clients, so it's a doddle to operate, and with *kdemultimedia* installed it also offers an embedded media player. Impressive stuff. Bandwidth restrictions, auto search-retry and a handful of other aspects can be tweaked, and though the

supplied documentation is shallow it's all very simple to use.

While it may not be at version 1.0 yet, from our testing we can safely say that *Apollon* is fully usable for day-to-day file-sharing jobs. The range of supported protocols via *gift* is superb (as with *Gaim*, it's always good to have one consistent front-end for different networks). We'd love to see *Apollon* adopted as KDE's main P2P client, enriching the desktop suite even further. Very highly recommended.



The attractive and easy-to-use interface is one of *Apollon's* many good points.

HTML EDITOR

Peacock

■ VERSION 1.9.1 ■ WEB <http://peacock.sourceforge.net>

With ISPs and portal sites such as Yahoo! GeoCities providing stacks of free webspace, just about every regular Web user has a page or two up for viewing. Linux is quite healthy in terms of site creation software at present – *Quanta*, *Bluefish*, *Mozilla Composer* and others all assist in building HTML pages, and now *Peacock* jumps in for some action. The original version was built around GTK 1, and this new release has made the leap to GNOME and GTK 2.

Unless you have a full GNOME 2.4 installation with all the necessary development libraries and headers, building *Peacock* from source can be a bit of a messy chore. We've supplied a .deb and RPM for Red Hat 9 on our coverdisc; if you're using the latter, you may need to update a few packages in

order to get things working properly. We managed to get *Peacock* going by mutilating a stock Red Hat 9 setup with various bits from *Rawhide*, but building with a proper GNOME 2.4 is the easiest way.

Peacock sports a familiar word processor-esque front end, with a tab bar for multiple document editing and toolbar for general HTML components. By using *GtkHTML 3*, the rendering is very much WYSIWYG and as a bonus doesn't have the full weight of *Gecko* behind it. Together with multiple undo/redo steps, tables, images and a basic template system, *Peacock's* generated HTML (4.0) is uncluttered and easy to modify by hand. Right now, the Preferences box hasn't been implemented, nor has the spelling checker, but otherwise the program



Peacock in action on a page with the table-construction box popped up.

appears to be in reliable shape. A 'file properties' box for editing the title and other HEAD tags is a sorely missing feature as well.

Hardcore HTMLers will probably find *Peacock* somewhat lacking in comparison to *Quanta*, *Bluefish et al*, but here the focus is different, and

anyone wanting to make a basic site with little HTML experience should give it a try. It's easy to approach, generally a fast performer and we had no major stability troubles; once the other planned extras are fully implemented, *Peacock* will be a splendid companion to the browser-only *Epiphany*.

WIKI SERVER

Yawk

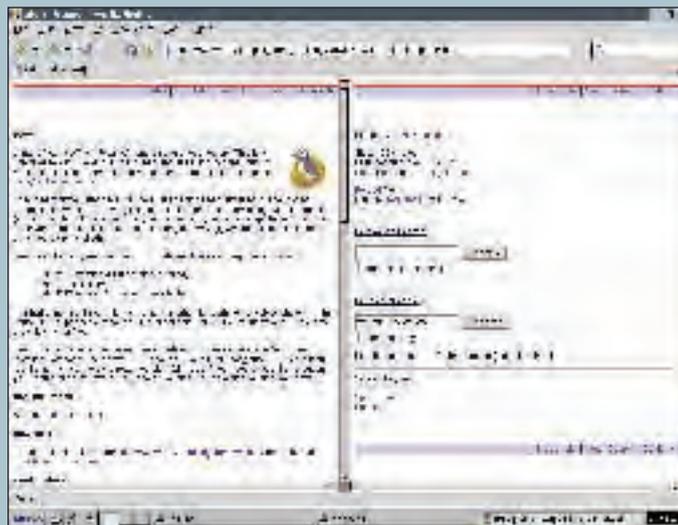
■ VERSION 1.21 ■ WEB www.awk-scripting.de/cgi/wiki.cgi/yawk/

Wiki (the name derived from the Hawaiian word 'wiki wiki', meaning 'quickly') is a system where multiple users can edit and manage web site content from any browser, anywhere. It's sometimes described as a democratic – and even anarchic – publishing system, and it's taking off for all kinds of projects. *Awk*, meanwhile, is primarily an interpreted text processing language with many variants in use today. Combine both, and *Yawk* is the end result – modestly named reader 'unixgod' suggested that we have a look, and right he was too.

In co-operation with a webserver, *Yawk* provides a small and speedy Wiki suite through its *Awk* codebase. You'll need at least version 3.1.0 of *Gawk* (*GNU Awk*) installed for the

script to run correctly, and all recent distros should have few troubles here. No specific webserver is required; however, the configuration docs are geared towards *Apache* so that's what we would recommend. The *INSTALL* file details the few steps of copying and editing, and it takes only a minute or two to get started (*parserdir* has been replaced by *bindir* in *wiki.cgi* but the documentation hasn't updated that, so watch out).

Yawk uses CSS for its visual style, and by default it supplies a clear and minimal layout sans fancy frills and distractions. A sweet touch is the ability to split the browser window into frames for editing in one and browsing in the other. The formatting tags system is fairly easy to become familiar with (though the use of '___'



Even *Yawk's* own site is constructed in *Yawk*! Right pane showing file info.

for italics and '%%' for underline seems back-to-front), and sports tables, images and nested lists. Excellent stuff.

A few glitches need sorting out even though *Yawk* is long past a 1.0 release – an unfinished bold marker in a document makes the bottom navigation bar bold too, for example, but otherwise it runs fluently.

Searching facilities are available, as is an info page for file sizes and modifications, and while it's not brimming with extras it's very pleasant to work with and easy to modify. If you're looking for a fast and bother-free system for online document collaboration on a fairly small-scale site, *Yawk* should definitely be investigated.

WORD PROCESSOR AND DESKTOP TOOLS

Flwriter

■ VERSION 0.14 ■ WEB www.oksid.ch/flwriter/index.html

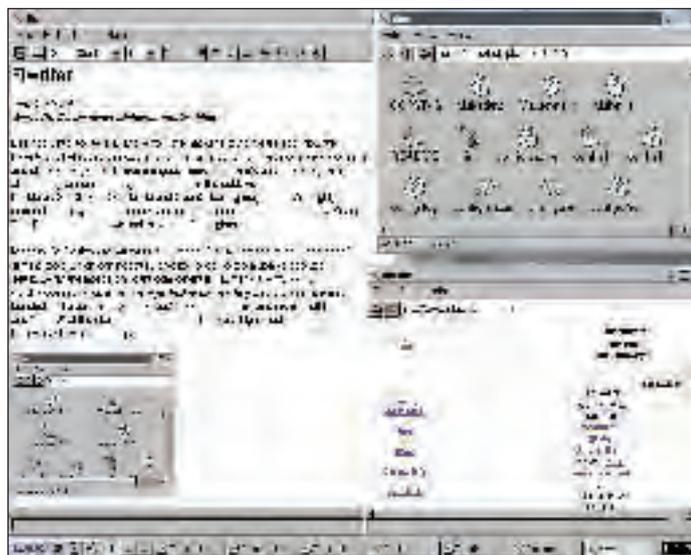
Desktop Linux today is a far cry from the hacker's toy of years ago. Now we have two fully-fledged desktop environments, a bunch of marvellous web browsers and office applications galore. These are good times, but there's only one snag with all this progression: older systems are increasingly being left behind in the cold. Powerful and flexible as GNOME, KDE, *Mozilla* and *OpenOffice.org* may be, they're highly demanding on system resources and need a modern box to run smoothly. Reader Tom Evans suggested we look at *Xd640* for a lighter alternative.

Meaning 'X Desktop at 640x480 in 16 megs', *Xd640* intends to deliver a set of slimline tools which can hopefully give low-spec boxes a usable desktop setup. By far the most significant component is *Flwriter* – the

WYSIWYG word processor – but a file manager, text editor and tiny web browser are also bundled in. To build from source, you'll need to compile and install *xutf8* and *ftk-utf8* at the start; the latter is a modified *Fast Light Toolkit* library used by all the progs.

Flwriter's front-end is ultra-simplistic and minimal, with a single toolbar for common operations. Basic features including Cut/Paste, Find/Replace and zooming are supported, along with image insertion (PNG only), an onscreen keyboard for foreign characters and RTF export. XHTML with embedded images is used as the default file format, so the resulting files are viewable in a browser, and on the whole it's a likeable little tool for basic documents.

The web browser, *Flspider*, is too rudimentary for regular use (*Dillo*



The word processor, file manager and web browser busy at work.

would be a better choice), while the limited file manager mixes GNOME and KDE icons and is acceptable for quick filesystem navigation. A spreadsheet is in the works too for inclusion in a future release. Realistically, though, *Xd640* doesn't strive to be an all-round power-user's desktop and the lack of features isn't a problem – Red

Hat won't be Bluecurving this any time soon. But for bringing a 486 laptop back to life, or giving an old box to a child for learning, *Xd640* is a cheerful suite of mini-utils.

For those interested in licensing issues, there's some very interesting links to discussions about GPL vs BSD licenses on the website.

BOMBERMAN CLONE

BomberClone

■ **VERSION** 0.10.1 ■ **WEB** www.bomberclone.de/core.html



Gadzooks! At least that Blue Guy is firmly stuck in Crisisville.

Bombs, then. It's strange how the metaphorical usage of the word differs depending on which side of the Atlantic you are situated. In the US, if

something is 'da bomb', that means that something is great and superb; yet in the UK, for a movie, song or piece of software to 'bomb' means

that things are going badly, and that none of the original investors are going to break even, let alone turn a profit. Yet in the computer game sphere, bombs have always been strangely appealing and game designers have enjoyed making them an integral part of the proceedings. *Bomberman*, Hudson Soft's top-notch long-running franchise has spawned an uncountable number of variations and clones thanks to its 'simple to play but ages to master' fundamentals. In *LXF12*'s *Game Roundup*, we checked out *Clanbomber*, and now *BomberClone* has joined the fray.

Thanks to SDL, *BomberClone* runs under Windows as well as the usual UNIX flavours. Pleasingly, no esoteric libraries are required so it should compile on every recent distro without difficulties. A binary tarball is also available (and on the coverdisc), but you'll need *glibc 2.3* to run it, and unlike earlier releases, all data files are folded into the main archive so nothing else is necessary.

The overall goal of *Bomberman*-style games is this: blow up the opposition. You're thrown into a 2D maze with human and/or CPU rivals, and with the help of bombs you can

blast through walls and set traps. Power-ups are available to increase the number of bombs laid simultaneously, double your running speed or widen the explosion radius. It's simple, silly and life-ruiningly addictive. Even with all the fast-paced action there's room for strategies – leading a player into a dead-end and executing the final blow is outrageously satisfying.

BomberClone offers the usual single player mode with a configurable number of CPU-driven players, and the AI behind them is suitably challenging (although you'll notice that they occasionally get stuck!) A multiplayer mode for Internet play is also available; but you'll need a fast connection to avoid dodgy latency (and messy oblivion).

Aside from the headache-inducing colour schemes of some of the maps, the visuals are mostly smart with twee rendered character sprites. Keyboard redefining appears unimplemented in this release and the sound subsystem needs a bit more work too, but it's still hugely playable. With a bit more spit-shine and other tweaks, *BomberClone* will rival the originals as an addictive, time-eating little animal.

3D BREAKOUT-STYLE GAME

Briquolo

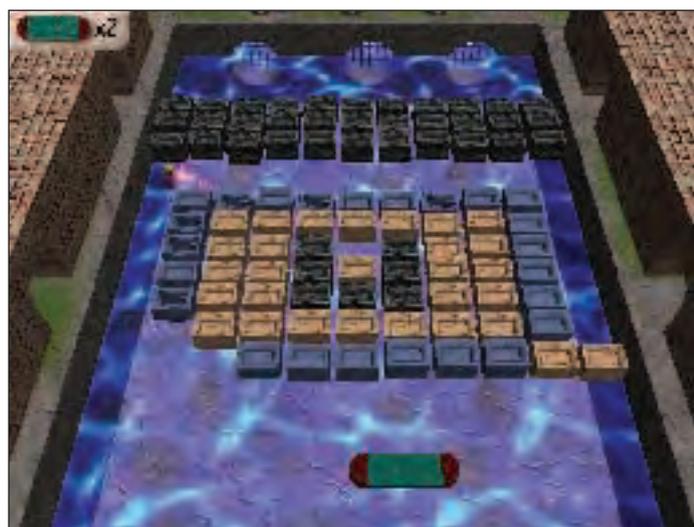
■ **VERSION** 0.4 ■ **WEB** <http://briquolo.free.fr/en/index.html>

The user manuals for modern computers, if at all in paper form, tend to be flimsy brochures with little substance and dollops of self-promotional hyperbole. Go back to the Sinclair ZX Spectrum though, and its chunky handbook contained all manner of useful tidbits – the +2 book even includes a full *Bustout* game in BASIC. Many an hour was spent by kids typing it in line-by-line and, amazingly, *Breakout* clones still live on. Whether you're feeling twinges of nostalgia for the likes of *Arkanoid* or just crave simple, raw, pure playability, *Briquolo*, a 3D OpenGL interpretation, could provide the solution.

It's always satisfying to see games that don't require a googolplex of obscure library dependencies, and *Briquolo* instantly scores points here: just SDL and related libs are needed

for compilation. No binaries are available at present, but the usual quick **./configure, make and make install** (as root) will build the game without difficulties (you'll need the *SDL_devel* and *SDL_mixer-devel* or equivalent packages installed too).

It's immediately apparent that much care has been put into *Briquolo*'s appearance; the silky colour-cycling title screen, anti-aliased fonts and magnificent in-game graphics all give a good first impression. The gameplay adheres to the traditional *Breakout* design – keep a ball in play by moving a paddle and obliterating blocks on the screen while you're at it. Sounds a breeze? With a complex layout and the ball flying everywhere, lightning-quick responses are essential, and real experts can even apply a bit of side to the ball, bar-room pool-style...



Once the ball starts bouncing up there, a few seconds of rest are in order.

Peppered with explosions and lighting effects, *Briquolo*'s gorgeous 3D visuals thankfully don't detract from the game itself, and purists can opt for an old-fashioned overhead view. The 16 supplied levels are designed with competence, and the level editor is a superb addition and creating zany arenas is a great deal of fun, however good you are at the game. All it really

needs to top-off the cosmetic side of the game is a soundtrack – something to keep the adrenaline pumping without being an annoyance.

Unsophisticated as the game idea is, *Breakout* clones are enjoyable as little desktop asides and few come more jazzed-up than *Briquolo*. It's attractive, customizable through the level editor and very well executed. **LXF**

Mandrake 9.2



If you're going to make the most of our new distro exclusive, you're going to need to install it! **Neil Bothwick** gives us some important pointers...

Mandrake has just released its newest Linux distribution: 9.2. These two pages will concentrate on what is on the *LXF* coverdiscs and how to install it; there is an alternative guide to installing it on page 106. The DVD comes with the

full download edition of Mandrake Linux, four CDs-worth, installable directly from the DVD. When we included Mandrake 9.1 with *LXF41*, the build process unfortunately omitted a couple of important packages when squeezing the distro down to two CDs, so we have allocated some more space for it this month and included another 100+MB of packages on the third CD.

Installing MDK 9.2

Mandrake is probably the easiest Linux distro for a new user to install – even easier than installing Windows on most systems! Even if you bought your PC with Windows pre-installed, you shouldn't have any difficulties installing Mandrake Linux to run alongside it in as dual-booting system, just follow the instructions here and the prompts given by the installation process. Having told you how easy it is to set up, don't run away with the idea that Mandrake is in some way simplistic, it is equally suitable for 'power users' too (whoever they may be), and the installer allows you to control the

choices it makes if you feel confident doing so.

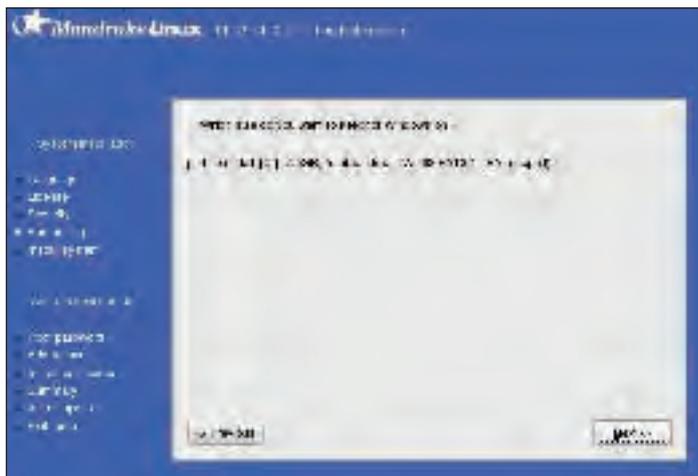
The easiest way to install is by booting from the DVD or first CD. If this is not possible, there are several alternate ways of starting the installation. These are covered on the CD pages. Note that the screenshots are from an install of Mandrake 9.2 RC2, the final release candidate. This is because these pages were written before Mandrake was able to send us the final 9.2 release for the discs.

The initial splash screen allows you to type some options before starting the install. Press **F1** to see the choice. The most likely use for this is to change the resolution of the install screens. Type **vgalo** or **vgahi** to change from the default 800x600 screens to 640x480 or 1024x768 respectively. Typing **text** gives a text mode install. You may need this if you have less than 128MB of memory in your computer. The graphical installer uses a lot of memory, because it loads everything into a ramdisk. Running text mode doesn't mean you won't have a GUI on the installed system, you can still install X from text mode.

The next few screens ask for basic questions to confirm your hardware and installation details. If you have an existing Mandrake installation, you will be given the choice of a new install, which will completely overwrite your previous setup, or an upgrade. Choosing upgrade will preserve most of your previous settings, as well as any other software you had installed.

Alongside Windows

The installer will resize a Windows partition to make room for Linux, but there are a couple of steps you should take in Windows before you start. First, run **chkdsk c:** from a command prompt to ensure your Windows filesystem is error-free. Then defrag the drive. This will ensure that the resizing process works as well as possible.



When the installer offers to resize the Windows partition, you may not see the drag knob on the slider, as in the screenshot. If that happens, just click on the area you can see to bring it into sight. Or it may just be a quirk of the *VMWare* setup that I used to test this.

PARTITIONING YOUR HARD DISK

Preparing for your Linux install

A hard disk is divided into partitions, each of which is accessed as a separate data storage area. A standard windows setup has a single partition occupying the whole drive, so you may not have even noticed that it was partitioned before, but adding a second operating system will require that each has its own portion of the disk. Unlike Windows, Linux installations generally use two or more partitions, the exact number and layout depending on what the machine is being used for.

Ask ten Linux users how you should partition your hard disk and you will probably get ten different answers. The

simplest setup has a swap partition and a root partition. The swap partition should be roughly twice the size of the computer's RAM, but probably not more than 512MB. It is used for temporary storage of data to free up memory. The root partition (/) contains everything else.

You can also have separate partitions for other sections of the filesystem hierarchy, /usr, /var/, /home and /usr/local are the most common candidates for this. However, this normally only complicates the situation. It is generally worth having /home on its own partition. This contains your own data and configuration files,

making it a separate partition means you won't lose it if you reinstall. It also means you can share one /home partition between more than one distribution, if you like to experiment with the different distributions on our cover discs.

Mandrake's installer uses *DiskDrake*, which will resize existing partitions for you, even Windows NTFS ones. It also makes reasonable choices about how to best use the available space. By all means use the custom partitioning option to make your own choices, but if unsure, feel confident about the automated choices.

ADDING SOFTWARE SOURCES

The contents of the install discs are added to Mandrake's list of known packages during installation; however, many more are available. If you installed from the CDs, there are the rest of the packages on the DVD, also on Mandrake's download mirrors. Then there are the contrib packages and the *Penguin Liberation Front*. The latter

contains many packages that aren't included in Mandrake for various reasons, such as the questionable legal status of *libdvdcss*, needed to play commercial DVDs. Adding extra sources is dead easy, point your browser at <http://plf.zarb.org/~nanardon/index.php> and follow the instructions that you are presented with.

See the box below left for details on partitioning your hard disk. Unless you have specific needs, it is generally best to accept the defaults given by the installer. Then it is time to choose which packages you want installed. In contrast to some other distros, Mandrake's installer doesn't provide options like 'laptop configuration' or 'install everything'; so you need to select which groups of packages you wish to install. You may be tempted to select everything on the basis that you may need it later. Resist this temptation! There's no point in installing software you don't need, it increases the disk space and time taken by the installation, adds potential security risks and makes backing up your system later a far more time-consuming task. Mandrake's Software Manager makes it child's play to add anything you may need later, so only select the groups you need.

There is also an option to select individual packages on the next screen if you want to fine-tune the installation. That screen also has an option to save your package selection to an MS-DOS formatted floppy disk. This is useful if you need to reinstall at a later date, it saves trying to remember what you chose last time.

You should select the Development section, or at least select *gcc* and *kernel-source* in the individual package section, otherwise you will not be able to install programs from source. If your computer is on a network, you should select the Network computer server group, in order to install *Samba* for file and printer sharing.

Press the Install button and the process starts. This can take a while, depending on the speed of your PC and how many packages you chose, so it may be time to put the kettle on. If you're installing from DVD, hard disk or over a network, you can leave it to its own devices (groan), CD users will need to change discs when requested.

When the package installation is finished, you will be asked for a root password. It is vital that you do set this (personally, I believe that someone at Mandrake should be very ashamed for including a No Password option). The root account should only be used for administrative tasks, such as installing software, and you will be prompted when it is needed. With no password, you could easily have root access, which can be dangerous, without realising it. It also makes your computer more susceptible to attack when online. Then you create the user account that you will normally use.

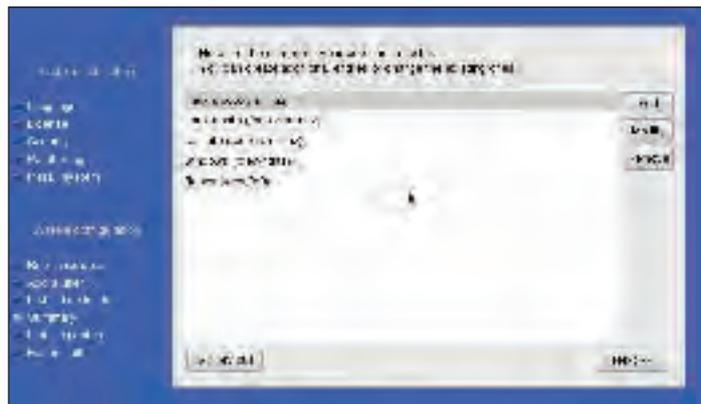
The bootloader is what starts up Linux, or gives you a choice of operating systems, when you boot up. If you have no other operating systems installed, or only Windows, take the MBR option. This will pick up your Windows system and give you a menu to choose which to use when booting. It is advisable to create a boot floppy here too. You may need this if you ever reinstall windows, as the Windows installer will overwrite your bootloader, making Linux inaccessible. In that situation, boot from the floppy and reset the bootloader in *Mandrake Control Center*.

Depending on your hardware, Mandrake's installer may or may not pre-configure the graphical interface. If it doesn't, you can do this from the Summary screen, where you can also change to any other configuration option. Select your monitor, graphics card and resolution here, and make sure you test it before committing to it.

The final step is to check for any updates that have been released since these discs were created. It is a good idea to do this during the installation. Not only does it check that your Internet connection is working and ensure your system is up-to-date, it also adds the relevant source for updates to Mandrake's software manager. This makes checking for updates later that much simpler.



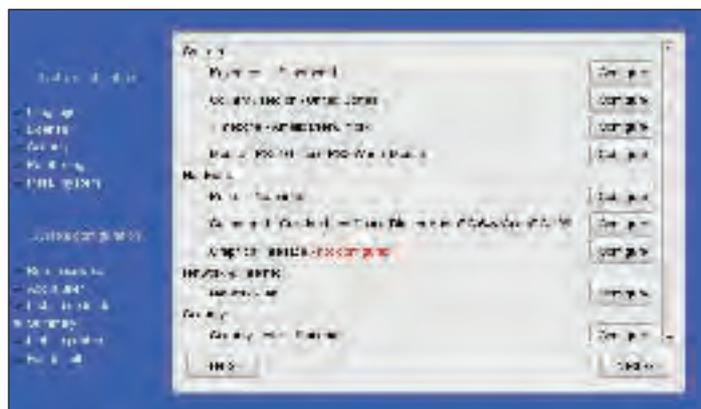
The individual package selection screen gives you a great deal of detailed information on each package.



The installer will detect an existing Windows setup and create a suitable bootloader menu.



Make sure you test your chosen resolution settings before you proceed.



Check everything is configured in the Summary screen before moving on.

« Using the KDE desktop

Kicker is flexible and can be positioned on your screen wherever you need it.



The default desktop environment in Mandrake is KDE. While there are many bits of software behind the scenes that make Linux work, this is the most noticeable, because it's the desktop environment that controls the way you run and use most of the applications. Many of these applications are provided with KDE too – everything from email clients and image viewers to PIMs, games and office software. Over

the next few pages, we are just going to cover the main parts of KDE you should get to know.

Konqueror is a dual-purpose piece of software which acts as a file manager as well as a web browser. With multi-paned and tabbed display areas, this very flexible tool is a neat way to work with files and documents wherever they may be.

The Kicker Panel

Kicker is the name given to the toolbar strip along the bottom of the KDE screen. Actually, although this is the default setting, you can actually set the kicker bar to appear on any edge of the screen, resize it or only have it appear when the pointer enters a particular part of the screen. The

kicker contains the KDE menu, shortcut icons for various applications, the virtual desktop switcher, representations of running software and, to the right, system tools such as battery level indicator, a clock, the *Korganiser* alarm icon and more if you have them configured.

The *KMenu* is designed to make it easy to launch any of your other desktop applications. A pop-up menu like that shown below with your applications organised into categories will appear when you click on the 'K' cog icon. Simply select the application you wish to run from the menu and off you go. This is also where you can log out of your current user session and log in as someone else or indeed shut down the computer.

EXPLORING THE KDE DESKTOP

Key parts of the new desktop explained

DRIVES

Removable drives will appear as shortcuts on the desktop. Double-click to open them

KONQUEROR – FILE MANAGER

View and manage documents and other files in Konqueror. You can customise the listing mode to suit your needs

DOCUMENTS

Individual files often have previews and file associations with standard Linux applications

TRASH CAN

Put files in here when you don't want them anymore, and remember to empty it!

KONQUEROR – BROWSER

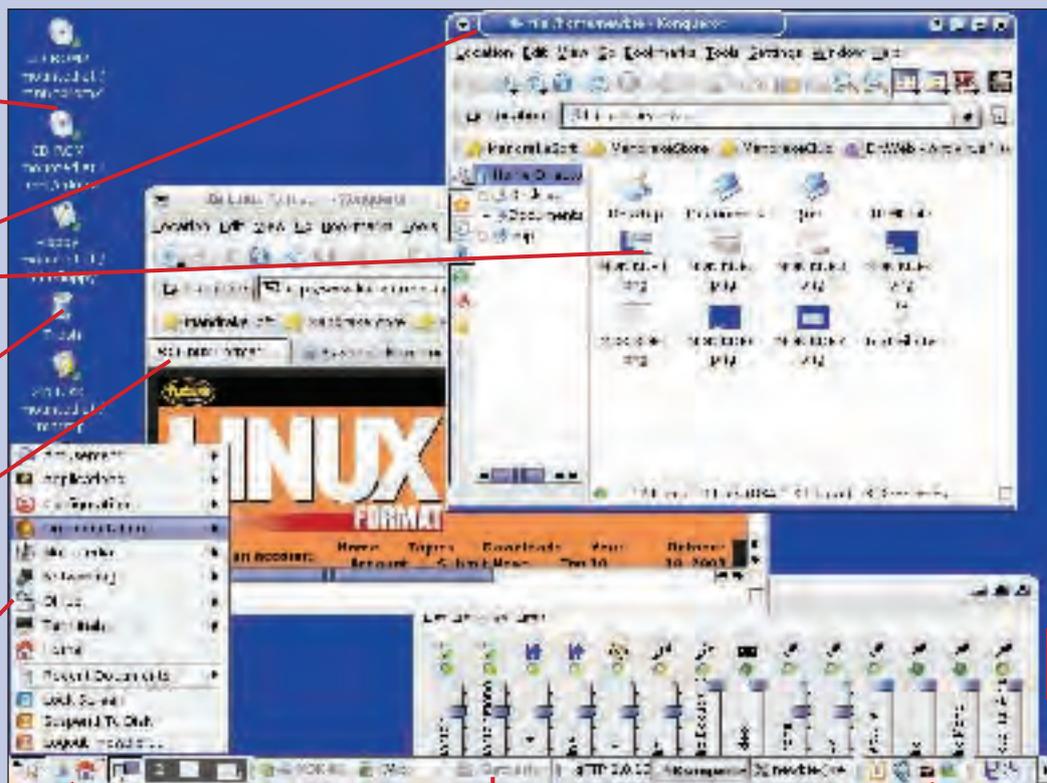
Konqueror also works as a web browser, with tabbed views and everything else you might want

THE KMENU

This is the 'Start' menu for KDE and includes most of the installed applications by default

SHORTCUTS

These are quick links to common apps – drag items from the *KMenu* here to customise it



VIRTUAL DESKTOPS

Switch between desktop views here – each desktop can have its own unique personality!

RUNNING PROCESSES

All running apps appear here, even if minimised – click them to bring them to the front of the desktop

SYSTEM TOOLS

Links to common system tools like Audio controls, the clipboard and others are available here

This version of Mandrake includes a new feature – ‘Suspend to disk’. When activated, instead of closing down active applications, the OS actually dumps all the relevant data and contents of memory to files on the hard disk. On startup, the exact system state can be restored. This is an experimental feature, and your experiences with it may vary, depending on the hardware you are using. You should also be aware that this could be a potential security risk if someone else may have access to your computer.

The virtual desktops mean that you can leave application windows open on one desktop, then switch to another for a clear screen to run other applications on. In practice, you might want to leave a console open on one desktop, office stuff on another, maybe a web browser on the third and so on. Each desktop can be customised to a certain extent, including changing the wallpaper to make it easier to remember where you are. It's possible to give these desktops names instead of numbers if that makes things easier.

As with most of KDE, virtually every aspect of the *Kicker* is fully customisable. Try right-clicking on some of the icons for example, and you'll get a pop-up menu allowing you to move the icon, change the application it links to, or even change the graphic itself. More icons can be added or removed as you wish, and pretty much every item on the *Kicker* can be repositioned.

As far as the *Kmenu* goes, more-or-less the same rules apply, but you will find it easier to use the *MenuDrake* configuration utility to adjust the properties here, which will then also be mirrored in the GNOME desktop menus if you later decide to use that as an alternative.

Konqueror

You will probably find the KDE app you use the most is *Konqueror*. This excellent software is both a file manager and a web browser. With split screen and tabbed viewing options, it has more features than you are likely to discover at one sitting.

When in file manager mode, *Konqueror* can be used to automatically launch the correct application for any number of different file types. These

CUSTOMISING KDE

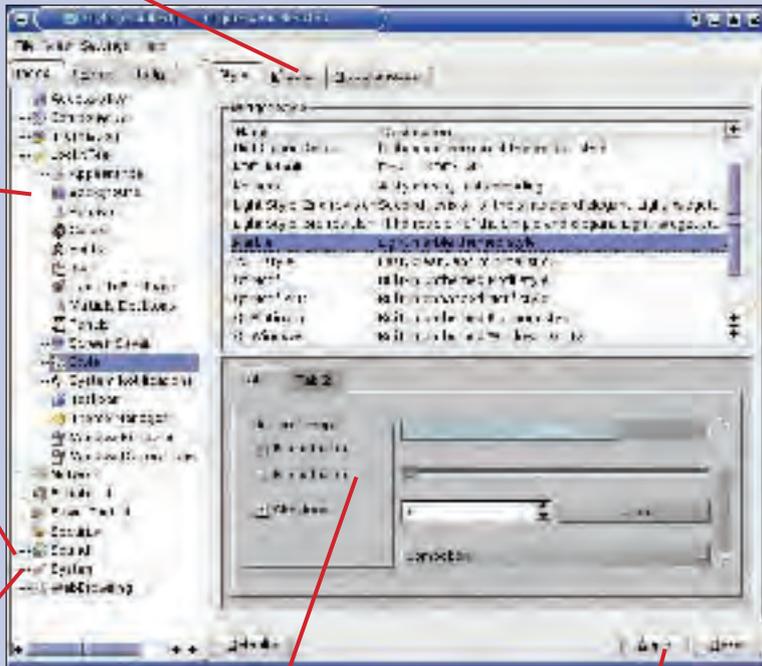
Using the *KDE Control Center* to configure the desktop

Most configuration tools have several tabbed panes for different settings

Choose a section from the tree. 'LookNFeel' is where all the eye-candy options live!

The sound section also includes default settings for Vorbis and MP3 files as well as CD audio preferences

Check out 'System' for some more advanced options, including the kernel configuration tool, boot manager and more (experts only!)



Most of the configuration tools will give you some sort of preview option

Make sure you hit 'Apply' to save the changes

are completely configurable, so if you don't like the default settings, simply edit the preferences.

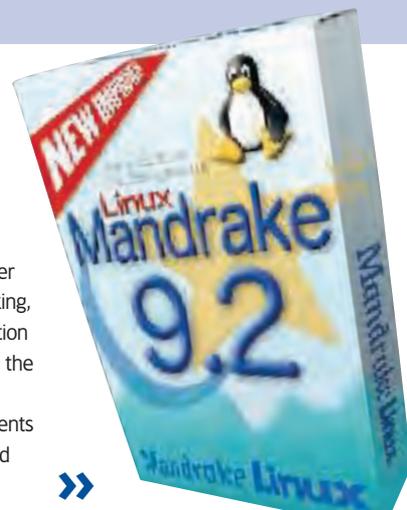
As a web browser, *Konqueror* has improved in leaps and bounds. You can also use it for ftp and network browsing. If you want to get the most out of this versatile software, choose the '*Konqueror* introduction' from the Help menu.

KDE Control Center

The *Control Center* configuration tool is used to set the default behaviour for the whole desktop environment. It can also be a useful tool for getting information about the currently running system. Run the *Control Center* from the *Kicker* or the *Kmenu* to change these settings. The *Control Center* splash screen will display basic information, while the panel on the left displays a tree of configuration options like in the annotated screenshot above.

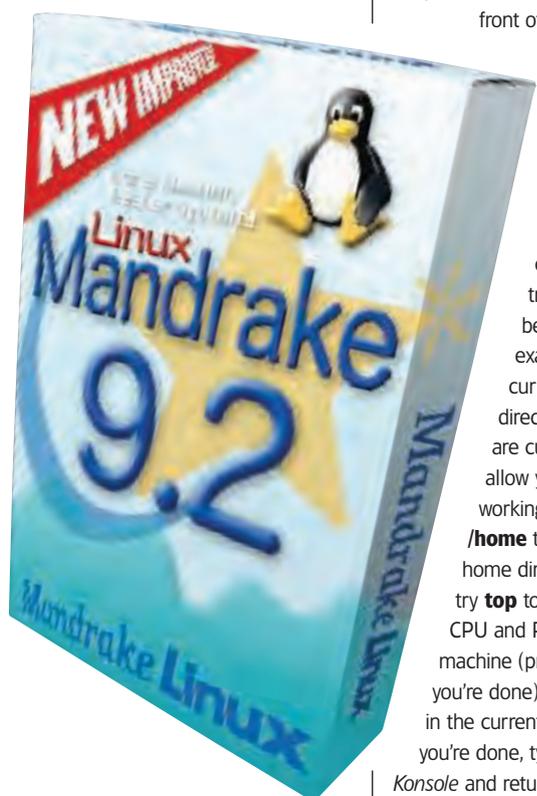
Choose the 'Look and Feel' section to alter the physical appearance of the desktop, including the overall theme,

wallpapers, icon sets and size and plenty more. This is the area of the *Control Center* that most new users want to explore first as it gives them the opportunity to personalise their computing environment. The other subsections are also worth checking out though. The Accessibility section for instance, can be used to alter the keyboard configuration, while the System section includes components for adjusting the boot settings and even configuring the Kernel.



Kmix is an easy way of configuring all your audio options in one place.

MANDRAKE 9.2



Controlling the console

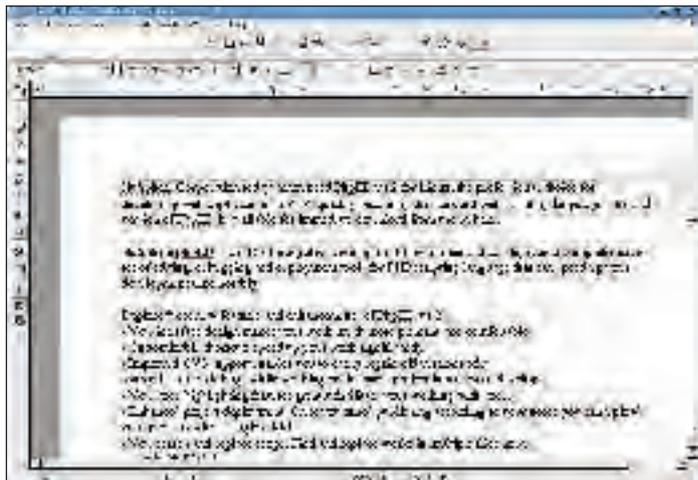
Although Mandrake has a very friendly front-end, power users will almost certainly want to delve into the world of the command-line at some point, and this is accomplished by using the application *Konsole*. The *Konsole* icon is a picture of a monitor with a shell in front of it, because it's often

known as a shell window – see, there is some logic there!

Once you're in *Konsole*, there are a number of commands you can try out to find your bearings. **pwd**, for example, will print the current working directory, or where you are currently, and **cd** will allow you to change the working directory – try **cd /home** to change to the home directory. You can also try **top** to print out current CPU and RAM usage on your machine (press Q to quit when you're done), or **ls** to list the files in the current directory. Once you're done, type **exit** to quit

Konsole and return back to KDE. Note that applications that you launch from the *Konsole* will also be killed when you close your *Konsole* window. If you're very new to Mandrake, you probably won't need to use *Konsole* all that much if at all – it's a very powerful way to interact with your computer, but you can work your way up to it slowly!

OpenOffice.org is central to MS-compatible productivity in just about every distro, and Mandrake boasts the latest stable version.



Word processing & spreadsheets

Mandrake comes with two power-packed office suites at your disposal, *OpenOffice.org* (*OOo*) and *KOffice*, each of which have their own advantages. While *KOffice* is much smaller and faster, *OOo* has far superior import and export filters for *Microsoft Office* – you can share documents almost seamlessly between the two packages.

To get started with *OOo Writer*, the word processor, click the K button, select 'Office', 'Word processors', then 'OpenOffice.org Writer'. The *Writer* interface is quite different to that of *Microsoft Word* and so will take a little getting used to. Luckily the help files are very comprehensive, and should be able to guide you safely through your early steps.

Unlike *Microsoft Office*, the programs that make up *OOo* are very closely integrated with each other, which means when you select 'File' and 'New' from *Writer*, you can start a new spreadsheet, a new presentation, etc – *OOo* will automatically start the correct program for you. *Writer* will have most if not all of the features you're used to with *Word*, including a macro recorder, embeddable charts, and automatic highlighting of spelling errors. This advanced functionality is equally reflected in *OpenOffice.org Calc*, its *Excel*-like app, and *Presenter*, its *Powerpoint*-like app.

By default, *OOo* saves in its own internal format – if you're not intending to share files with non-*OOo* users, this is the best format to use. However, it can also save documents in the equivalent *Office* format (*Word* format for *Writer*, *Excel* format for *Calc*, etc).

If you don't need the full-blown complexity of *OpenOffice.org*, *KOffice* has a lot to offer also. It too can import and export using *Microsoft Office* filters, but they aren't as good as *OOo*'s filters. The advantages to using *KOffice* is that it is much better integrated with the rest of KDE, and also that each *KOffice* app loads a great deal faster – very often you load a document in *KOffice*, print it off, then close the application before *OOo* has even finished loading! Here at LXF, we tend to use *KOffice* in much the same way as Windows or Mac users use *Notepad* – jottings and text-only files – and save the capable *OpenOffice.org* for more intensive word-processing tasks.

Playing games

Although many people think that Linux is only good for servers, Mandrake is working hard to dispel that myth by providing an excellent selection of games along with the distro. Perhaps the most popular game bundled with Mandrake is *Frozen Bubble*, which is a thoroughly cutesy arcade-style game that is a cross between *Tetris* and *Breakout* which has featured several times on the LXF coverdiscs. The object of the game is that you have to fire coloured balls up in the air so that they land on balls of the same colour. When you have three or more balls together, they will all disappear, scoring you points. Once you've mastered basic aiming, you should try bouncing the bubbles off the left and right walls, as this will help you get to otherwise unreachable places. Launch *Frozen Bubble* by clicking the K button, selecting Amusement, then Arcade, and finally *Frozen Bubble*.

If you've configured your graphics card properly by downloading and installing the latest drivers from your manufacturer, you'll likely find the most enjoyable games are *Tux Racer*, *Cannon Smash*, and *ArmageTron* – each of which require a fairly good video card on your behalf, as they are all 3D games. Furthermore, configuring your graphics card to use 3D usually takes quite a bit of work, and you should read the instructions carefully.

Listen to music

You can satisfy all your music needs with Mandrake through *XMMS* (the *X Multimedia System*). This program might seem a little eccentric at first, but it's actually modelled after the popular *WinAmp* MP3 player for Windows – you can even share skins and many other plugins between the two. To start *XMMS*, click the K button, select 'Multimedia', 'Sound', then '*XMMS*'. Click the up-facing arrow to open a music file for playing – *XMMS* handles MP3 files, Ogg Vorbis, and indeed most other music file formats.

To change the way *XMMS* looks ('skinning' is a popular sport in many parts of the Linux community!), and also access most of the other *XMMS* features, right-click the application to bring up the main menu. Skin changing is available under 'Options';

'Skin browser' – *XMMS* comes with a small collection of skins for you to try out, but there are hundreds more available for you to download for free on the Web. Once you have your sound sorted out, your next stop should be the visualisations – where *XMMS* will create a variety of animations that move to the rhythm of the current track that is being played. A basic selection of visualisations comes with *XMMS*, but, as with skins, you can of course get many more online. Some of the more advanced visualisations will require you to have downloaded and installed the correct graphics driver for your system. To quit *XMMS*, just right-click and select 'Exit'.

Manipulating graphics

Whether you're a budding artist or an accomplished neo-modernist, you're in luck – Mandrake comes an excellent image manipulation suite in the form of *The GIMP*. To start *GIMP*, click the K button, select 'Multimedia', then 'Video', and finally '*The GIMP*'. The interface will take a little getting used to to begin with – almost of the suite's functionality is reached by right-clicking on the image you're working with and going from there. To get started with a new image, click 'New..' from the File menu and select the size of the picture you want.



***XMMS* is so similar to *WinAmp* you can use skins to fool Windows users!**

The default toolkit in *The GIMP* is modelled around *Adobe Photoshop*, so you should feel at home if you're from that environment. For example, to drop text into your picture, select the italicised **T** symbol from the toolkit and click on your image to type. You can access *GIMP*'s filters by right-clicking your image and selecting the Filters menu. There's also the 'Script-Fu' menu that contains many user-written filters that are often combinations of the standard filters designed to produce a given effect.

To save an image once you're done with it, right-click a picture and select 'File', 'Save As..' – most file operations are performed from this same menu. michael J Hammel started his popular *GIMP* tutorial series back in *LXF38* – order your backissues on page 99.

Managing your finances

Mandrake comes with *GnuCash* to handle your home or small business accounting needs, and you can start it by clicking the K button, selecting

'Applications', 'Finances', and finally '*GnuCash*'. If you've ever used Intuit's *Quicken* on Windows, *GnuCash* should feel more than vaguely familiar – you need to add information about your costs each month, and also your income, and it can help you analyse it all. If your bank supports the creation of QIF files (and most banks do), you can import that directly into *GnuCash* and save yourself all the hard work. *GnuCash* includes a variety of wizards (called 'druids', just for variety) to help you accomplish various simple tasks such as adding a new account, and these can help you learn the intricacies of the program quickly.

Perhaps the best feature of the program is to allow you to program in scheduled transactions such as standing orders or automatic debits, and it will use that data to predict your cashflow in the near future. Also be sure to check out *GnuCash*'s powerful charting tools to help you visualise your profit and loss

If you get stuck, be sure to have a look through *GnuCash*'s help files – they are very comprehensive, and should help you solve any problems you have. Alternatively, *LXF43* carried an in-depth tutorial covering all that you need to know to start getting to grips with *GnuCash*.

Burning a CD

Mandrake includes many applications for burning CDs, most of which we have covered in our roundup back in *LXF43*. All have their place, depending on the type of project you are trying to create, but probably the best in terms of integration with the Mandrake desktop is *K3b*.

To start burning a data disc for example, all that you need to do is select files or directories in *Konqueror*, right-click and select 'Create Data CD with *K3b*' from the pop-up menu. This launches the burning software and automatically creates a new project containing the files you selected. If you are running the software for the first time, you may need to visit the Settings menu, but *K3b* is actually pretty good at scanning for and recognising devices.

Slip a blank CD into the drive and press the burn icon in the toolbar, and you'll have a shiny, data-filled CDROM in no time.



Mandrake 9.2 spoils you for choice with applications for burning CDs.



***The GIMP* comes with Mandrake and can handle most graphical tasks from simple photo retouching to professional image creation.**



« Configuring Mandrake with *DrakConf*



One of the many reasons why people have turned to Mandrake for their choice of distribution has been its ease-of-use. Although this simplicity can be tested in everyday use of applications and software, a core test is how easily your system can be configured with it.

Mandrake makes use of the *DrakConf* tool, which is more commonly known as the *Mandrake Control Center*, which we touched upon a few pages back. With the 9.2 release of this tool, few areas cannot be configured graphically using one of the many subsections of the application.

Here, we will take a more in-depth look at some of the most commonly used parts of the *Control Center*, and try to determine how easy configuring your system really is.

Software Installation

Software installation and updates are essential for a fully working system, and the Software Management part of the *Control Center* configures this. The first things to set up are your software sources with the Software Media Manager; the tool will allow you to specify network sources for packages and CD/Local sources. With this complete, you can now use *RpmDrake* to select packages to install. Within *RpmDrake* you will see a list of categories and you can browse the categories for software. You can then select a package from the list and it will be installed when you click on the Install button.

If you would like to install RPMs that you have downloaded from the Internet, you can download them to your hard disk and double click on icon to install the RPM. You will be

notified of any package dependencies (packages required by the RPM) that you may require.

Configuring 3D cards

Graphics card (and 3D acceleration) is configured with *XfDrake*. Your card will have probably been detected during the installation of Mandrake 9.2, but using *XfDrake* you can fine tune your resolution/colour depth and test that your new settings work.

Internet Connection

The *DrakConnect* tool can be used to select network hardware and get connected to the Internet. If you click on the Launch the wizard button your network hardware will be detected, or you can select your device from a list. With the hardware installed you can then configure the connection by clicking on the 'Configure Internet

ANATOMY OF DRAKCONF

Tailor your Mandrake system to meet your exact needs

BOOT

Here you can configure how your system starts

HARDWARE

You can configure your different hardware devices here and how they operate

MOUNT POINTS

This section lets you configure how you can access disks and other media

NETWORK AND INTERNET

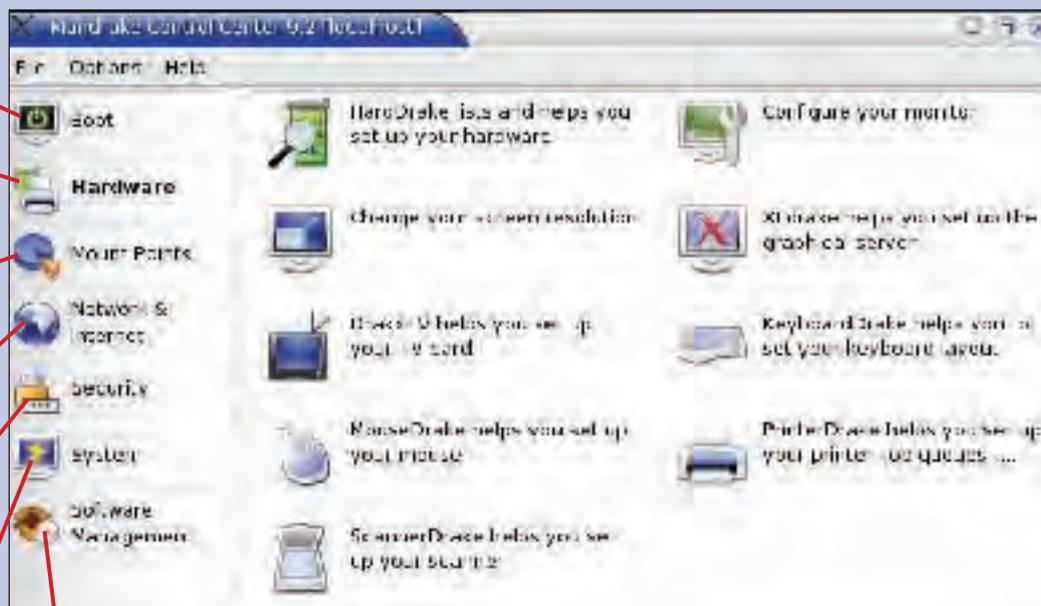
You can set up your LAN, ISDN, Wireless and Modem access in this section

SECURITY

You can easily configure a firewall and set how strong security is on your system

SYSTEM

This section allows you to configure exactly how your system operates, and setting up system fonts, backups, services, users and more



SOFTWARE MANAGEMENT
This section will let you add/remove new software and download essential security and application updates

« This area is used to select items to configure, and when you select an item, » its configuration program will be loaded.

Access' button. *DrakConf* will now take some information from you about the connection (such as dial-up phone number, username/password etc) and you can then activate the connection by clicking on the 'Connect' button.

Adding Users

User login accounts can be added to your computer by clicking on the *UserDrake* icon in the System section of the *Control Center*. If you click on the Add User button, you can enter some information about a new user (such as username/password) and the account will be created. The default checkboxes in the box will set up a sensible user account.

You can allow lots of users to share the same group by first adding a group with the 'Add Group' button, and then clicking the 'Edit' button and its 'Groups' tab to select the group for the user.

Printers & Scanners

Printers can be configured with the *PrinterDrake* tool in the hardware section of the *Control Center*.

Mandrake will try to auto-detect your printer, and if it is not found, you can select the printer from a list.

Additional printers can be configured using the 'Add Printer' button. If you need to access a printer on a network, you can use the 'Configure CUPS' buttons to set up remote printers. Printers connected to a remote computer can be configured with the 'Printer Sharing' button.

Scanners can be configured with the *ScannerDrake* tool also in the Hardware section of the Control Center. Mandrake will attempt to configure your scanner automatically, but can be manually selected by using the 'Add Scanner' button. If you would like to share the scanner on a network, you can click the 'Scanner Sharing' button and select the checkbox to share the scanner. You can then click the button

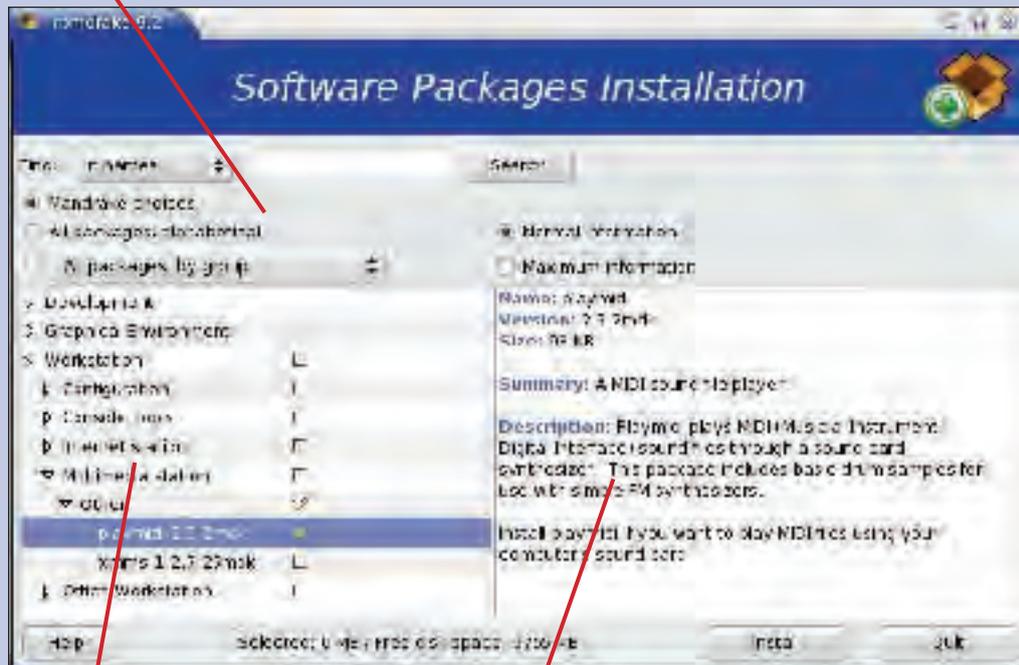
LXF SPECIAL OFFER! 15% off joining MandrakeClub

Whatever level of Mandrake Linux user you are, there is exclusive content that will appeal to you that can only be had by joining MandrakeClub. Turn to page 107 in order to take advantage of LXF's exclusive membership fee discount!

ANATOMY OF RPMDRAKE

Installing your software packages

This area lets you search for a package and select how the packages are listed



This area displays the packages available for installation and each has a checkbox so can select that package to be installed

Here information is displayed about the package that is selected in the left-hand pane

to configure the hosts allowed to use the scanner select the 'All Remote Machines' option.

Firewall

Firewalls are essential to be secure on the Internet, and the *DrakFirewall* tool in the Security section can be used to configure a simple firewall for your machine. Using the tool you can select which software and services should be allowed Internet access, and all other connections will be blocked. It is important to only select the software you actually need, otherwise you may be vulnerable to an attack.

Accessing disks

Accessing disk drives is called 'mounting' in Linux, and can be configured with the tools in the 'Mount points' section. In Linux you need to specify a directory on your disk where a drive can be accessed (such as /mnt/floppy for a floppy drive). The tools within this section allow you to create these mount points easily and

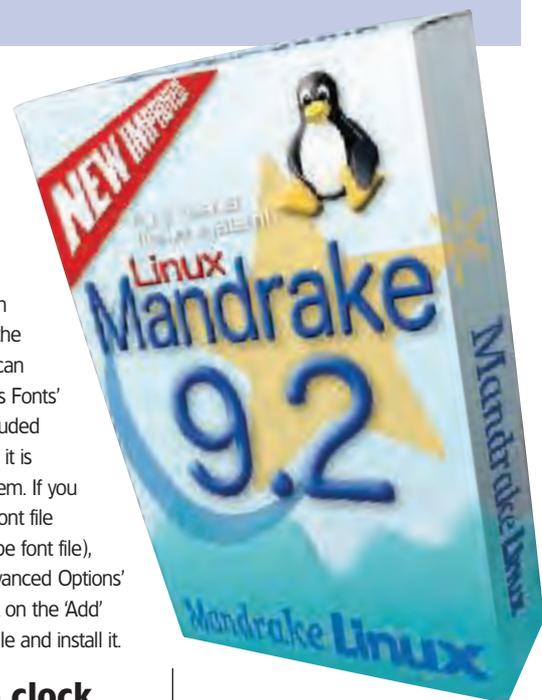
when you access the directories, the disk will be automatically accessed.

Add fonts

This can be done with the *DrakFont* tool in the System section. You can use the 'Get Windows Fonts' button to get the included fonts with Windows if it is installed on your system. If you have downloaded a font file (such as a .ttf TrueType font file), you can click the 'Advanced Options' button and then click on the 'Add' button to select the file and install it.

Setting the clock

You can set the clock by clicking the button in the System section. Then use the scroll boxes to select the time and click on the date to select the date. Your timezone can also be selected by clicking on the 'Timezone' button. [LXF](#)



What on Earth is... CVS?

Tracking changes in source code would be tricky without it, says **Jono Bacon**.

» **Most distros seem to come with CVS, but I've never been sure what it actually is or does – please enlighten me!**

Well, CVS is a source control system. This is a chunk of software that is used to save changes and modifications to source code.

» **Right, I'm not sure I understood much of what you just said. Could you elaborate at all?**

Sure. Let's look at this with an example. Let us assume that you are writing a program and you create your first version. Let us now assume that a few weeks later you are editing the code, and realise

you have made a fundamental mistake and would like to revert to the code you wrote a few weeks previously. In this situation you are up the proverbial excretory creek without a paddle.

» **Let me guess, CVS could have helped here.**

Spot on. CVS allows you to save your changes to a special place called a repository after every development session. When you save these changes (known as committing) you also add a log message to say what you did, and CVS will give the addition a version starting at 1.1. Each new addition increases the version (1.1, 1.2, 1.3 etc) and CVS merges your additions to the code.

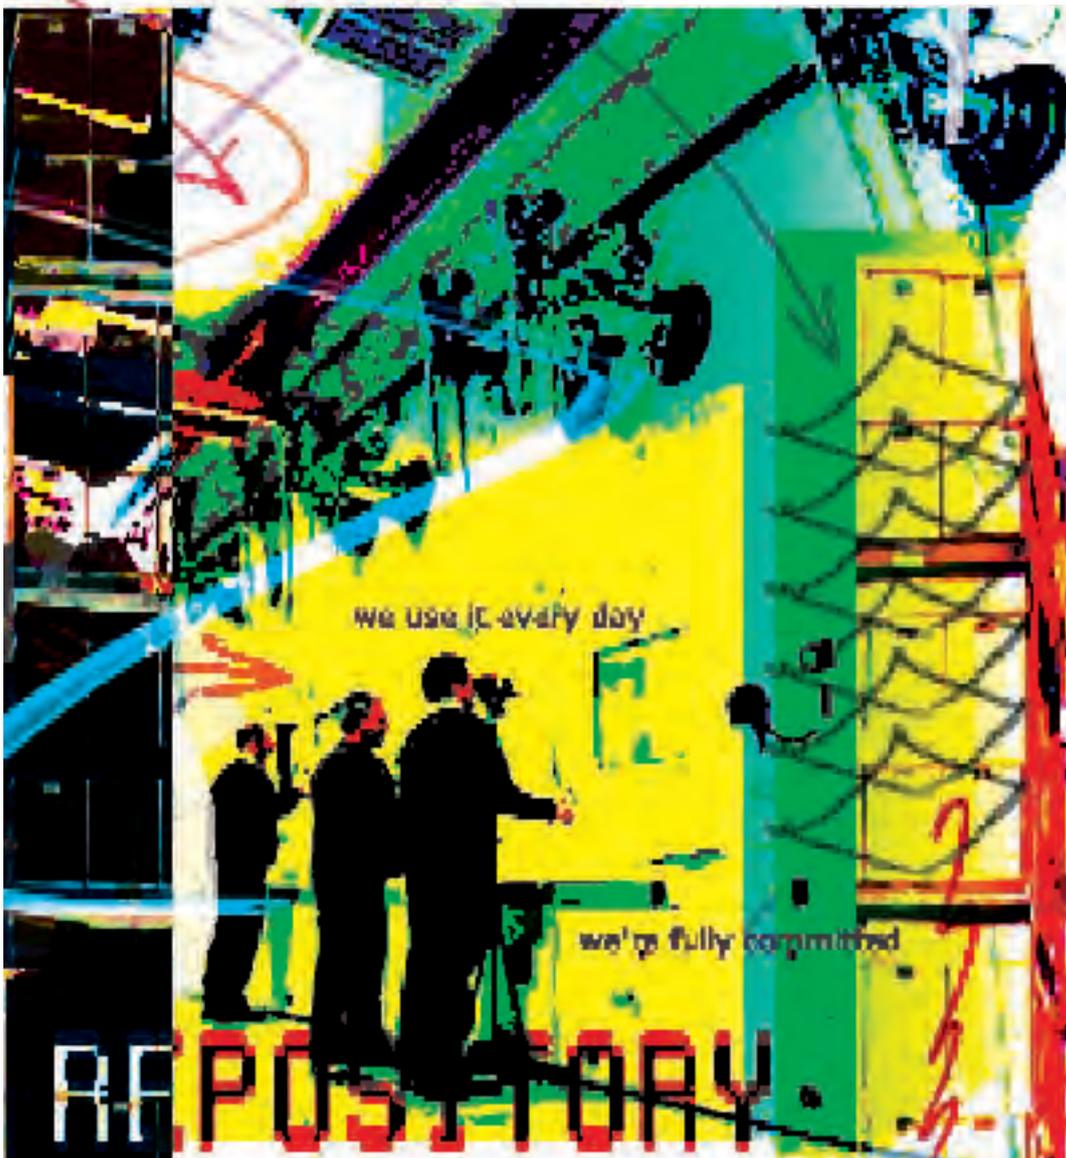
» **What do you mean by "merging additions"?**

Well, if CVS was to save the entire file each time, the system would not be very efficient would it? To increase the all important efficiency factor, CVS only adds the changes you made to a file – this also helps you see what you added in the new version.

» **Is that all CVS does then? Not exactly exciting is it...**

Hang on there, we've only just started. CVS is also a popular system for enabling distributed development – the way in which most Open Source and Free Software projects are written. This method of development allows multiple developers from all over the world to work on the same project. Naturally, all of these developers need to work from the same code, and CVS is used to manage this.

Many projects in the Linux world make use of CVS such as KDE, GNOME, XFree86, *Blender* and more. This is one of the main reasons why CVS is so popular – it is used practically by hundreds of developers every day.



»» So how does CVS handle all of these different developers?

CVS makes use of a security system where a user needs to be added to the CVS server. This user will then have access to commit changes to the code. You add a user to the CVS repository's password file and you can even add anonymous access so people can download from the server but not commit changes to it. Many projects offer this facility so you can get the latest cutting-edge version of their software.

»» OK, so I can see how this may be useful, but do you really need to be able to revert back to old versions of the code?

Although this feature is used every so often by most developers, it is there as a failsafe. Ask yourself the same question about backups – they are only needed rarely, so are they worth making?

»» Ah, good point. So developers only go back to older versions when something goes wrong – surely CVS must mess the code up when adding the changes each time though?

Generally CVS is pretty good at keeping the code in its original form. Sometimes there may be a conflict, but CVS will tell you when this happens so you can commit them again.

»» Hang on a second, surely there must be conflicts if two developers are trying to commit code at exactly the same time?

CVS usually places a lock on the file in question that could be modified at the same time. When the first developer has finished committing, the lock is removed and the second developer can commit. Although CVS is a good system for source control, it is not perfect however, and there are indeed conflicts here and there. CVS cannot be used to fix bad management practices!

»» What do you mean bad management?

Let us assume you were running a software project and you had three developers. You would generally ask each developer to work on different parts of the source code to ensure there are no conflicts. You would discuss with the developers what needs doing and who will do what. If this does not exist, things will break with CVS or no CVS.

»» So can I only use CVS for writing C/C++ code or something?

Of course not, CVS will work with virtually any type of media that can be patched (media in which you can add/remove parts of it in plain text). This includes virtually any kind of source code, documentation, web pages and other things. You

can also import binary files (they cannot be patched) and when there is a new version, CVS will just import the entire new version. As an example, I use CVS myself for all of my writing and web development work.

»» Web development?

Yes, in fact, this is one of the most useful areas I have found for CVS. I develop a website and add it to my CVS server. When the site is ready to go live, I log into the webserver and check the site out. This means that I can move my entire website to a new server with a single command. This is also good for running the same site on multiple servers – you just check out the same CVS module.

»» I see; so CVS is used for updating software as well as storing it?

That's right. CVS is specifically a source control system, but it is also a great way of moving software around, and updating it with just the changes; it is cases like the web development example where CVS is really useful. Another example would be keeping your configuration files in a CVS module. If you then set up a new machine, just check out the module and you are configured.

»» OK, you have persuaded me. So how do I use it? You said earlier that I could get the latest bleeding edge software from Linux projects. How would I go about doing this?

Well, the first thing is to choose which project you want to compile. As an example we will use the 3D modeller *Blender*. If you go to the homepage www.blender.org there is a page called CVS.

On this page there are instructions on how to connect to the CVS server. Most projects include a page such as this, and connecting basically involves logging into the machine with a long terminal command (as shown on the *Blender* page), or you can set the environmental variable CVSROOT to point to the repository. For *Blender* the CVSROOT is:

```
:pserver:anonymous@cvs.blender.org:/cvsroot/bf-blender
```

You can set this in the bash shell with:

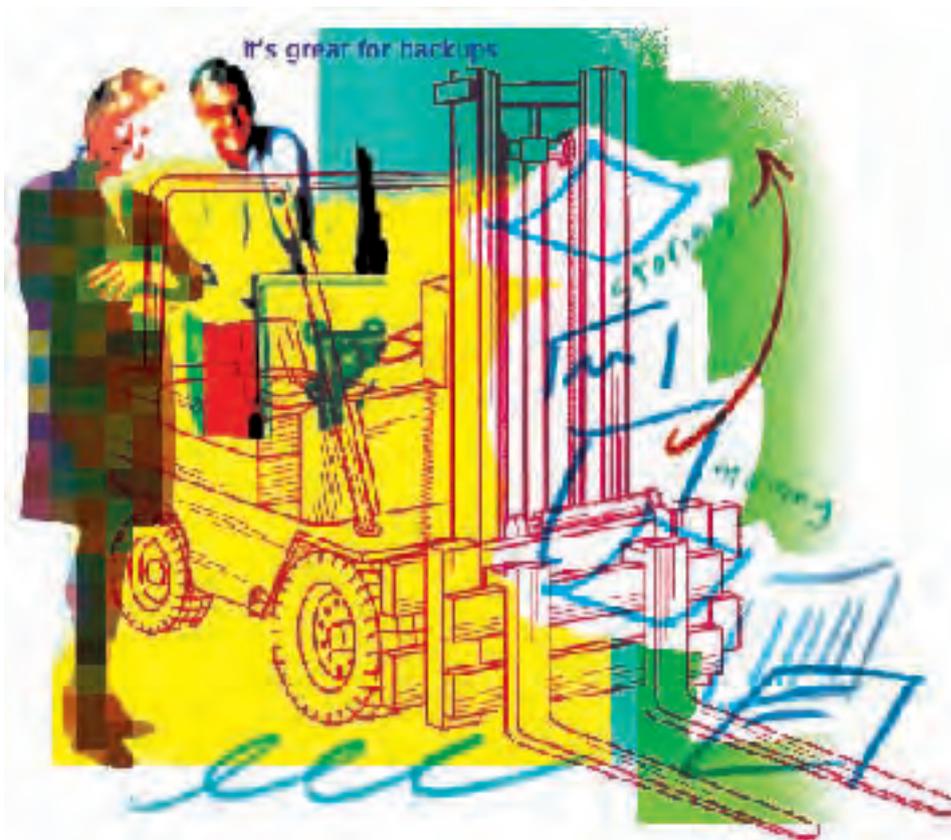
```
export
CVSROOT=:pserver:anonymous@cvs.blender.org:/
cvsroot/bf-blender
```

»» Whoa! Hang on there a moment, what does that CVSROOT thing mean? Please explain!

The CVSROOT comprises of three parts, each separated by a colon (:). The first part refers to how users are authenticated. This server uses :pserver, which means a password server. There are other forms of authentication including SSH if required. The second part specifies the user and host where the CVS server is located. In this case it is the anonymous user at cvs.blender.org. The final



WHAT ON EARTH CVS



part refers to the location on the CVS server where the repository is held; in this case `/cvsroot/bf-bender`.

You can then log in to the server with:

```
cvs login
```

For anonymous CVS access you generally just press enter when asked for a password. If you get an error that a `.cvspass` file does not exist, just create one in your home directory with:

```
touch .cvspass
```

Some versions of CVS don't like it when you don't have this file so you need to create it.

Right, I am logged in, so how do I get the *Blender* code?

To get the code, you need to know the name of the module in the CVS (a module is a directory that contains a project). This can often be found with the following command:

```
cvs co -c
```

This should list the modules in the CVS, but not *all* CVS admins set this up properly, and in the *Blender* case they haven't either. The *Blender* CVS page has a link to a page which shows the contents of the CVS however, and you can see on that page that the module is called **blender**. With this knowledge you can check the module out with:

```
cvs co blender
```

You'll then see a listing of files come from the server.

I see the files come out – the beginning of each line starts with a letter 'U'. What does this mean?

Whenever files come in and out of the repository,

each file has a letter at the start of the line to say what is happening to the file. The **U** in this case indicates that the file is new and is being checked out for the first time.

Other letters include **P** for when the file is getting patched (having code added/removed), and **?** refers to a file that is not in the CVS.

How can I look at the log entries for the files in the source code?

This is simple. You just enter the command:

```
cvs log file
```

You may want to put the output of the command into 'less', as many files have many log entries. So for a C code file you issue the following command to look at the log entries:

```
cvs log file.c | less
```

OK, now I have got some of the code, what do I need to do to update it later when the *Blender* hackers have modified the code?

Updating is simple, whether you have modified the code yourself, or are wanting the latest functionality provided by a development team. Just go into the *blender* directory and enter the command:

```
cvs update -PA
```

This command will check out all of the modifications since your last checkout and merge them into the code. The **-PA** options customise our updates slightly, and will remove directories that have been deleted, remove sticky tags and checkout to our directory.

OK, I get that, but what the heck are sticky tags?

Sticky tags and data are files and information that is associated with your checked out working copy that may not be relevant to the server copy. Although the feature exists, it is not absolutely necessary to using CVS, and you can refer to it in the manual. We will stick to the main features.

Fair enough. Right, let us jump in feet first and get down and dirty with CVS. How do I set up a beefcake CVS server?

Setting up a CVS server and going fully beefcake is not too difficult, but there are some important steps that we need to go through. The first thing you should do is actually get a copy of CVS. The main distribution of CVS comes with the client and the server, although they are often packaged separately for distributions. Check with your distribution first; it may well be already installed on your machine.

Where can I download it from the developers of CVS?

You can get it from <http://ccvs.cvshome.org/servlets/ProjectDownloadList>. Although there are binary releases from this site, this is mainly the source release which will mean that you have to compile it. That requires the usual `./configure`, `make`, `make install` jiggery-pokery. See the documentation at www.cvshome.com for specific instructions on this.

OK, so let us assume I have it installed, how do I set it up?

The first thing you will need to do is to create a CVS repository. This will be a directory that will hold the CVS files. Tradition seems to infer that the choice of directory is `/cvs`, although some use `/home/cvs` and others. It doesn't really matter, so create the directory. You then need to initialise the directory as a repository. This requires that you run a command on the directory (here we use `/cvs` as an example, change it for your directory):

```
cvs -d /cvs init
```

If all went well, a `CVSROOT` directory will be created in your repository.

So what on earth is the `CVSROOT` directory for?

This directory holds a bunch of administrative files that CVS uses to keep ticking over. We will look at some of them in a bit, so just be patient for now. We now need to set up our authentication system for users.

Aha... are we now going to use that `:pserver` jobby that you mentioned earlier?

That's the one. Like I said before, you can do this with SSH if you want for a more secure system, and

there is plenty of documentation online how to do this. Many people use the `:pserver` method however, so I will cover that here. To set up the `:pserver` (password server), CVS will authenticate users with a file in the `CVSROOT` directory called `passwd`. Before we can do this though, we need to get our CVS server running, and to do this we must add this to `/etc/services`:

```
cvspserver 2401/tcp
```

You must also add this to `/etc/inetd.conf`:

```
cvspserver stream tcp nowait root
```

```
/usr/local/bin/cvs/
```

```
cvs --allow-root=/cvs pserver
```

This will start the CVS server on port 2401 and you may need to poke a hole in your firewall so it is accessible if you need it to be. With this in place, we can then add some users to our CVS server. Adding users is done in the `passwd` file which is located in the `CVSROOT`. It is likely that this file does not exist, so you will have to create it. Each line has a different user with the following format:

```
[user]:[encryptedpassword]:[systemaccount]
```

An example of this could be:

```
bob:$df4fggtwedfdw334dfdf:cvs
```

In this example we have the user called **bob**, an encrypted password, and the system account we can use to perform CVS processes, which is optional.

»» How do I encrypt the password?

The easiest method of getting an encrypted password is to copy it from `/etc/passwd` or `/etc/shadow`. There are also some utilities to create encrypted passwords.

»» OK, so how do I add an anonymous user?

Well, you first need to add the anonymous user in the `passwd` file and leave the password part blank. You can then specify that the user account in the `/cvs/CVSROOT/readers` file. You are now pretty much set up.

»» How do I actually put things in the server?

This is all dependent on what you are putting in. There are basically three types of content you put in the repository. These are:

- directories full of code
- files
- modifications to files

To get us started, we will need to import our first module. Most people usually import a directory full of code that will constitute the module. Let us assume you are importing a directory called **myapp**. You would use this command:

```
cvs import myapp bob start
```

This command must be run from inside the `myapp` directory. The command says "cvs, I want to import the current directory into the repository and call the module **myapp** using my username which is **bob** as an identifier, and now start".

Whenever you add anything, the default editor will pop up and ask for a log entry. Enter a brief description of what you are adding, save and exit. You will then see the files added to the repository. If each file has the letter N at the start of the line, the file imported correctly.

If you would like to show the module name when someone enters **cvs co -c**, you must add it to the `/cvs/CVSROOT/modules` file. To do this, add the module name twice on the same line at the bottom of this file, like this:

```
myapp myapp
```

»» Are there any types of files I should not add?

The main files you should keep out of the repository are files such as configuration files and build scripts that have been generated that are not useful to someone compiling the code. CVS will attempt to ignore some of these files (such as `file.txt~` editor backup files), but it is best to remove them if possible. Also remove any hidden directories that have been generated (such as `.pics`).

»» OK, but how do I add legitimate files, OR REMOVE THEM?

Adding files simply involves entering **cvs add file** and then **cvs commit**. You then add a log entry and the file is added.

To remove a file, delete it first from your checked out copy and then enter:

```
cvs remove file
```

The file will still be in the history of the project, but will not be checked out in future copies.

»» Is there anything else I should know?



Well I think something you should be aware of is tools to make using CVS easier and more fluid. These include web-based repository viewers and graphical clients.

»» Web based what?

Web based repository viewers. These are web scripts that let you look at a repository, its modules, logs and other information and even request files that show the differences between two different source code files. One of the most popular web-based repository viewers is *ViewCVS*, and is available at <http://viewcvs.sourceforge.net>.

»» What graphical tools are available then and are they any good?

There are quite a few and many are good. Favourites include the KDE-based *Cervisia* (<http://cervisia.sourceforge.net/>), *LinCVS* (www.lincvs.org/), the GNOME-based *Pharmacy* (<http://pharmacy.sourceforge.net/>) and the java-based *jCVS* (www.jcvs.org/). All of these clients offer a means of graphically adding, removing, updating, committing, adding logs and other functions. There is also a Windows client called *WinCVS* (www.wincvs.org/).

»» Hang on a sec, did you see that? This is Linux Format magazine; not Windows Format!

Valid point, though unlike some of the more zealous Linux advocates, *LXF* supports the view that there is plenty of room for proprietary and Open Source/Free Software products to co-exist. This is where CVS is particularly handy – cross-platform source control. You can access your Linux-based CVS server and store your Linux, Unix, Windows, Mac and other projects on there. As such you may want to have a graphical client for each OS that you use.

Give CVS a go. I find it a genuinely useful tool that I use in every form of creative work for storing my code/documents etc. There are indeed alternatives available for you to try such as *Subversion* (<http://subversion.tigris.org/>) or *Bitkeeper* (www.bitkeeper.com/), but few of these have the userbase and popularity of CVS, and therefore not commonly used. CVS is by no means perfect, and *Subversion* is showing particular signs of interest for many developers, but CVS is still undisputedly king of the source control hill.

»» Well thanks...

One other thing. Remember that in these days of worms, viruses, crackers, hardware failures and other negative aspects of computer-using life, backups are essential, and CVS helps with this. If you keep all of your code/documents in a CVS repository, you can backup the entire server with a few commands. Combine this with a *cron* job that performs a **cvs** update every hour, and you have a complete backup server. [LXF](#)

Tutorials >>>

Our experts offer help and opinions on a whole host of Linux applications

YOUR GUIDE TO GETTING THINGS DONE!

Whether you are just starting out in Linux, or an experienced veteran, there's always more to learn. Every issue of *Linux Format* is packed full of practical advice, and nowhere is it more concentrated than in our tutorials pages.

Here you'll find expert guides to all sorts of things, from Basic Linux usage to understanding and deploying network solutions, from simple script coding to the complexities of Perl regular expressions, Java server apps and more. We aim to bring a good mix of tutorials to each issue, but if you have any suggestions for topics you'd like us to cover, why not contact us, by email at linuxformat@futurenet.co.uk or by snail mail, or log on to www.linuxformat.co.uk and post your suggestions in our special forums? Hope to hear from you soon!

Nick Veitch EDITOR

HOW CODE IS REPRESENTED

Including code in magazines can be tricky, but we hope our notation will help it become clear. When lines are too long for our columns, the remaining text appears on the next line in a solid blue box:

```
procedure
TfrmTextEditor.mniWordWrapClick
(Sender: TObject);
otherwise, there is usually a gap between lines:
begin
  mniWordWrap.Checked := false
end;
```

Usually, you'll find the code on our CD/DVD too.

THIS MONTH TEACH YOURSELF...

Beginners' tutorial

Linux home networking with either wireless or cables is easier than you think... **p66**

Compiler writing

Tidying up the code that we've written so far – loop incrementing and operator unification **p70**

Blender >>

On the eve of its return to Open Source, Jono has been monkeying around with the Game Engine **p74**

The GIMP

How to customise patterns, and you can even create your own clones of *Adobe Photoshop* brushes! **p78**

Practical PHP

Revolutionise your code by using sockets to build a PHP server that scans ports, accesses the web and more besides **p82**



ILLUSTRATION BY MARCUS STADE - www.blender3d.org

Server School

Introducing the core technologies and concepts behind firewalls and why you should use them **p86**

Scanning

Going mad trying to find drivers to make your flatbed work under Linux? SANE is the answer! **p90**

TIP OF THE YEAR! FREE WEB HOSTING FOR LIFE!

Mindfield Media is offering a lifetime supply of web hosting to anyone who gets a tattoo of one of its company mascots. Any person with skin is eligible for this unique opportunity!

THIS IS NOT A JOKE. Mindfield Media is an incorporated business located in Dallastown, Pennsylvania, offering affordable web hosting and design services to customers in the central Pennsylvania area since 1998.

"The tricky part has always been: How do you market an abstract service like web hosting?" comments Mindfield Media's CEO, Walt Miller.

"We've always found it incredibly funny that on every web hosting

company's website you'll see a picture of a server", continues graphic artist Brian Magar. "Every company looks the same. We thought we'd inject a little marketing fun into our web services; so far, we haven't seen anyone take it to the same level!"

Magar designed the look of three Mindfield Media mascots – a baby chicken (Hosty The Hosting Peep), a limbless robot (D.N.S. The Serverbot) and a bunny (Url The Rabbit) – that are gaining cult status with Mindfield Media's target market: artists, media, musicians and record labels.

"What does a baby chick have to do with web hosting? Absolutely

nothing, but that's kind of the beauty of it", says Miller. "We've had a great response. People send in pictures of our promotional stickers in use all the time. There's a nice word-of-mouth thing happening with this!"

Merely choose one mascot (or more!) from the cast of characters, read the rules of the promotion, download the appropriate template from www.mindfieldmedia.com/propaganda_tattoo.html and get your tattoo. And VERY IMPORTANT: Have someone take pictures as you are receiving that tat. They need solid proof that you are loyal enough to receive a lifetime of free web hosting!

NETWORKING BASICS

Beginners' Guide to Linux: Networking

Home networking is a big growth area as computers increasingly achieve a state of ubiquity – and it's not that difficult to set up, promises **Andy Channelle**.

We'll show you how to create a simple, peer-to-peer network with two computers which can be used to share files and an Internet connection. While we will try to keep things simple, networking can be a complex subject and may, occasionally, need a brief excursion in the world of The Command Line.

It's not uncommon these days for homes to own more than one computer; a laptop for mobile working, a PC for 'school work' in the kids' room, an old, obsolete Pentium cast aside in the rush to upgrade. And while a computer is great on its own, the benefits of networking them together really outweigh the difficulties you may encounter. MP3s can be accessible from any computer, one printer can serve the entire family, and everyone can surf the 'net at the same time.

The most simple network, a pair of PCs lashed together by a single wire, may be all you will need to do. Linux makes it quite easy, and the latest distributions from Mandrake, Red Hat and SuSE do much of the hard work for you.

A very simple network

INGREDIENTS

Two PCs (not surprisingly!) Both machines should have a reasonably modern Linux distribution on it (we will deal with connecting to Windows next time) with the network hardware recognised, though it doesn't matter if it is not yet fully configured. The big three distributions have decent graphical

tools, and we're assuming you're using one of these, rather than fiddling with .conf files.

Network cards these are the interfaces which allow your machines to talk to each other. Broadly, Network Interface Cards (NICs) fall into two categories: PCI for desktop machines and PCMCIA or PC Card for notebook computers. Most examples that you'll encounter are Linux friendly. You can also get USB NICs too, but if you're going down this route, make sure you check their compatibility with your distro vendor. Most have pretty good lists of compatible hardware on their websites.

CAT5 Crossover cable This is not the same as normal network cable that would usually go into an intermediary device such as a hub or switch. Make sure you get a crossover cable (Belkin colours them yellow), and that you have a long enough length of cable for your plans.

METHOD

Once installed, Linux will usually recognise most network cards at bootup and launch a 'new hardware wizard' to let you configure it. If this doesn't happen, you have to start the process manually through your distro's Control Panel.

While there may be a few differences in the way you enter information, the actual process of configuring a network card is broadly the same, so while we've used SUSE 8.2 for our examples here, you should be able to transpose the method to your chosen distribution.

Each network device is given its own number, so if the card is installed correctly, it should show up in the *Control Center* as device **eth0**. Highlight the device and select 'Edit'. We are dealing with two computers here with no connection to the outside world, so we need to opt for a static IP address and subnet mask.

In the IP Address space, enter **192.168.0.1** on Machine One and **192.168.0.2** on Machine Two. The subnet mask for both should be **255.255.255.0**. Exiting will then save the configuration and 'restart' network services so the devices are 'up' or active.

To test the connection, hook up the Crossover cable to both machines and on Machine Two open a terminal or Konsole and type: **ping 192.168.0.1 -c10**. This will send ten 64K packets of data from one machine to the other and monitor response times. The results should look similar to **Fig1**, above right.

If you get no response, check connections, make sure the LED lights on both network cards are lit (one should flash when information is sent or received) and ensure the configuration (IP address/subnet mask) is correct on both machines and that you are using a Crossover cable. If you still have a problem, run through your distribution's network card set up again, making sure the right drivers are automatically selected or manually chosen.

Once you have established a connection, you need to set up each machine to facilitate file sharing. This is done on a

SUSE's YaST provides a smart graphical front end for network configuration.



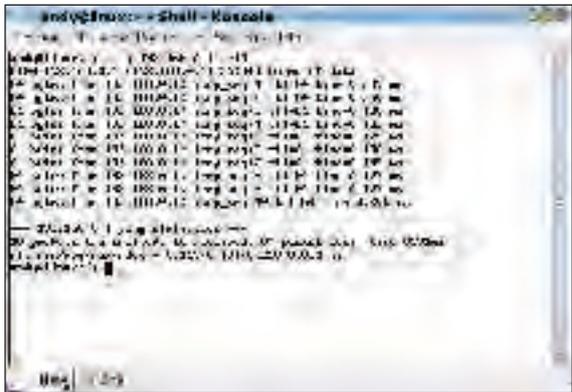


Fig1 Monitoring response times for ten 64K packets of data.

client/server basis, though each machine on the network can be both client and server should the need arise.

I got the music in me

For this real-world example, we'll be making a large selection of MP3s stored on Machine One available to the rest of the network which, at present, consists of a laptop. This is a two-stage procedure: first we need to define what is to be shared – the server element, and then we make it accessible from the remote device – the client element.

In the interests of command line literacy, we'll first examine how you can set up a shared directory in Mandrake using a terminal, before looking at a couple of graphical methods below. You'll need your Mandrake CDs handy.

TOOL BOXES

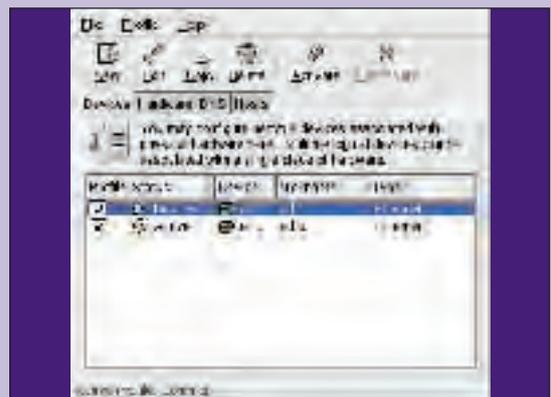
Configuring your network

It's worthwhile pointing out that the distros we have mentioned have two configuration systems. In addition to the KDE/GNOME *Control Center*, which is usually on the top level of the menu and – in the main – deals with individual users' settings, all three have their own method for 'system-wide' configuration.

SUSE *YaST2* (*Yet another Setup Tool*) is located at Kmenu>System>YaST2.

MANDRAKE *MDK Control Center* is accessible from the desktop in a default Mandrake installation.

RED HAT Red Hat breaks system configuration down into components, rather than having one place to do everything. Network configuration is at Kmenu>System Settings>Network.



Access to any of these tools requires your 'root' password.

- 1 Become root. Type **su** into a terminal/Konsole, hit return and enter your root password.
- 2 Install the correct software. In the terminal, enter **urpmi nfs-utils** and follow the instructions on the screen. This command simply uses Mandrake's package management tool to install *nfs-utils*. You can also select the package using RPM Drake.
- 3 Start up the portmapper. Type **service portmap start**.
- 4 Define directories to export. Launch *Kwrite* by typing **kdesu kwrite**. You will be prompted for the root password again before the application opens. Open the file `/etc/exports`. It will probably be blank, so, initially, we'll need to add a single line.

```
/home/andy/Music. 192.168.0.2 (ro)
```

This tells Linux which directory to export, the machine allowed to access it, and the ways it can be used. In this case we have a directory of MP3s (in `/home/andy/Music`) which we want to make available to Machine Two, configured with the address **192.168.0.2**, and we want to make sure that users of Machine Two can't change files, so we export the directory as read-only. To make the directory writeable, change **(ro)** to **(rw)**. Save the file.

- 5 Start the server. Type **service nfs start**.

On the client, you can now mount the shared directory. Create a directory in `/mnt` called 'DesktopPC' and, in a terminal, type:

```
mount 192.168.0.1:/home/andy/Music /mnt/DesktopPC
```

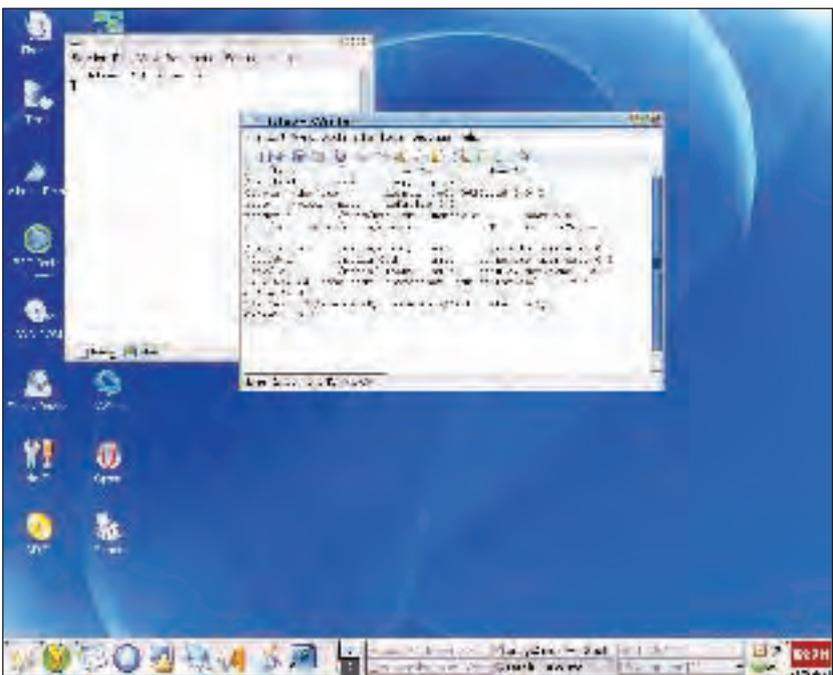
Substituting the paths for your own previously defined share. The first part of these entry defines the directory to be accessed and the second sets the 'mount point' on the laptop. In Konqueror, you can browse and play the MP3s via the `/mnt/DesktopPC` directory.

You can auto mount the shared directory by adding an entry in `/etc/fstab`. Open *KWrite* as root again (**kdesu kwrite**) and open `/etc/fstab`. The entry for an NFS share follows the same scheme as other entries in this file. For the above example, the line would be:

```
192.168.0.2:/home/andy/Music /mnt/DesktopPC nfs defaults 0 0
```

If this looks a little daunting, Mandrake's *Control Center* has a spiffy Internet and network setup tool which does the same thing with a graphical user interface, which you should find fairly straightforward to click through.

***fstab* is where you configure automounted devices. You'll need root privileges to edit it.**



TUTORIAL Beginners' Linux



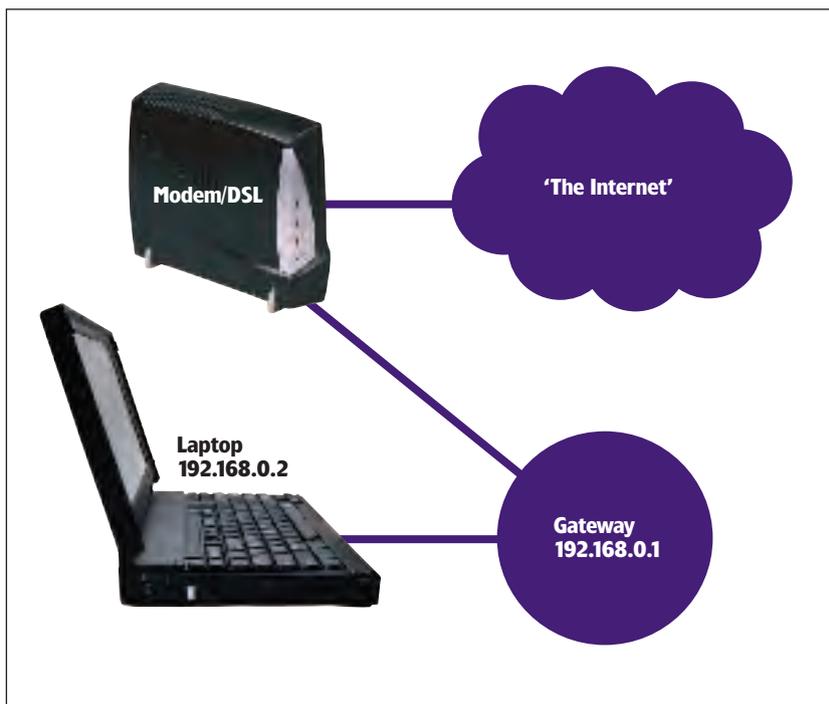
Konqueror makes file sharing easy, but should problems arise you need a bit of knowledge.

- ◀ KDE users have it even easier. To set up NFS shares on your LAN in KDE 3.1, you will first need to configure global file sharing.
- 1 Right-click on any folder and select 'Properties'.
 - 2 Choose the 'Local Net Sharing' tab.
 - 3 Click the Configure File Sharing button and check the option marked 'Allow users to share files from their HOME directory'. Hit 'OK'. You will need the root password to do this.

Now any user can right click on any folder in their directory, choose the 'Share' tab and select 'Shared'. This will allow access to the directory over both NFS and *Samba*.

Once properly set up, you should be able to navigate around the remote file system as though it were local, within the limitations previously set.

Sharing an Internet connection is the most common reason for networking home PCs.



Fishing for connections

While setting up NFS shares is a good idea, there may be occasions when you need access to a specific part of a remote system that hasn't been configured for file sharing, for instance if you want user access to Machine Two from Machine One and you're unable to get root privileges. In this case you can use *Konqueror's fish* protocol to securely connect to another machine and browse its file structure as though it were locally mounted.

To *Konqueror*, fish is just another protocol like ftp, http or smb, so using it simply involves typing **fish://** followed by the URL or IP address into the browser's address bar.

To access Machine Two from Machine One we do:

```
fish://192.168.0.2
```

The connection will be started and you'll be presented with a dialog box requesting a user name and password, which has to be valid for the machine you are connecting to. You could also do:

```
fish://root@192.168.0.2
```

to get root access to Machine Two if you had the password.

Fish uses *ssh* (secure shell) access and so doesn't need any server software running on the remote machine. It is also possible to connect using fish over the Internet, though as most broadband or dial up connections have a dynamic IP address, you'll need some way of divining the address of your PC. A number of solutions are available to do this, such as No-IP (www.noip.com), which features a Linux client for automatically updating your IP address and assigning a memorable URL to it. No-IP has a basic, free service available.

All join in

Probably the most common reason to network two computers is to share a single Internet connection, and while this may have been complicated in the past, distribution developers have made it almost foolproof, with a decent range of graphical tools to conceal the inner workings. For this example, we're using SUSE 8.2 but, again, Mandrake and Red Hat have similar tools available in the Internet and Network configuration panels.

As with file sharing, connection sharing is a two-stage process. First a few prerequisites: we will assume you have a working Internet connection on one machine and can ping between PCs with confidence.

Machine One has two connections, an external, Internet facing one and the LAN we configured earlier. Whether you connect to the 'net via a dial up, cable modem or DSL the process is roughly the same.

The external connection probably uses a dynamic IP address, meaning it is assigned to you as needed by the ISP's DHCP server. DHCP also takes care of configuring DNS (Domain Name System) addresses so when you type a URL into the address bar of *Mozilla* it can resolve it and access the correct website. We need to enable IP forwarding on Machine One and then configure Machine Two to use it as a gateway and access the correct DNS servers to resolve web addresses.

SUSE's network settings are accessed via KMenu>System>YaST2 under the Network Devices section, and you will need root privileges to access it. Choose 'Network card', highlight the device that links to your network and hit edit. We first need to activate IP forwarding, which is in the 'Routing' section. Make sure 'Enable IP Forwarding' is selected and hit the OK button.

Next click on 'Hostname and name server' (choose 'Accept' on the pop-up) and write down the IP addresses under the 'Name servers and domain search list'. These will be needed when

configuring the second machine. Click 'Back', 'Next' and 'Finish'.

On Machine Two, go into the Network Devices section, highlight the network card and hit 'Edit'. The first port of call, again, is the 'Routing' section, though this time make sure IP Forwarding is not enabled. Instead add the IP address of Machine One's second network card – in the preceding example, this would be **192.168.0.1** – to the Default Gateway bar. Hit 'OK'. Next go into 'Host name and name server' and insert the IP addresses taken from Machine One in the list. 'OK', 'Next' and 'Finish'.

Finally, open a web browser and type in an address...

If you're using a modem on Machine One, it should be set up with Dial on Demand enabled, so if you attempt to connect from Machine Two and the gateway is not connected, it will dial out.

The Network unleashed

Just a few years ago, home improvement TV shows would point out the silvery high-tech DIY projects which would see Mr. Toffee-Knose of Chelsea kitting out his designer loft with 500 meters of Cat5 cabling to 'future-proof' his investment. Basic wireless kit is now hardly more expensive than its tethered brethren, (D'oh!) but wireless offers the ability to work or play anywhere in the house, garden or close-by pub without the Cat5 umbilical cord and expensive wall replastering. And while most retailers don't shout about Linux compatibility (the fools?!), much of what you can buy will work perfectly well with a recent distribution.

Dabs.com, for instance, has a wireless PCMCIA NIC in its Dabs value range (basically a re-branded GWA-E06) for a ludicrous £18.99. A USB stick (part number EW-7117U) that does the same job for desktop machines is just £22.50. While the specifications on the company's website don't mention it, both devices are Linux compatible with drivers available on the Internet and, possibly, on the installation CD.

Wireless networking obviously introduces a couple of security concerns. The first is that neighbours and passing Warchalkers may want to borrow some of your bandwidth for their own surfing (which can often lead to you unwittingly breaking your terms and conditions on sharing connections). If you're simply using two wireless NICs to connect two machines, the simplest thing to do is choose a non-standard IP address (*ie* not **192.168.xxx.xxx**) for your gateway and then configure the laptop's default gateway with the same address. If you've gone wild and bought a wireless router, you can restrict access to specific MAC addresses, which are tied to network hardware, making it difficult to breach.

The other major headache is data security: sending your passwords, love letters and financial details floating through the ether unprotected is not a good idea. For this reason there is an encryption standard called Wired Equivalent Privacy (WEP) that is intended, as the name suggests, to offer a similar level of protection as a traditionally cabled network.

Unfortunately there isn't room here to go into configuring WEP, as much of the process depends on the devices you are using and your distribution, so a little judicious research with Google would probably be in order.

If you need more written information on the ins and outs of wireless networking, *802.11 Wireless Networks: The Definitive Guide*, by Matthew Gast (ISBN 0-5960-0183-5, O'Reilly) has a good selection of general and Linux-specific information. Alternatively, *Wireless Home Networking for Dummies* by Danny Briere, Pat Hurley and Walter Bruce (ISBN 0-7645-3910-8) takes a different approach but understandably is aimed more at Windows users – no slur intended. [LXF](#)

GLOSSARY

Network terminology

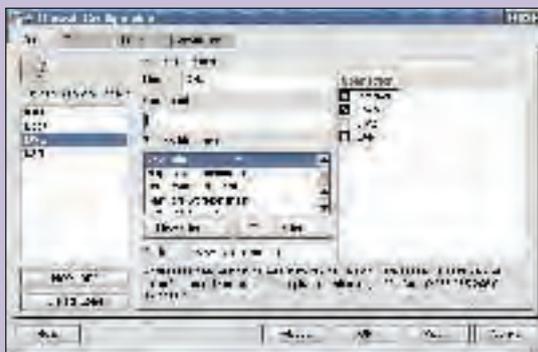
NFS Network File System. A method for mounting remote file systems on a machine. NFS drives are accessible as though they were local.

SAMBA A client/server architecture that allows you to integrate Linux/Unix and Windows computers on the same network. It is based on Microsoft's Server Message Block (SMB) protocol.

IP ADDRESS Every computer connected to a network – be it LAN, WAN or Internet – will have a unique IP address. This is the series of four numbers separated by a period (.) that, on the World Wide Web is obfuscated into a standard web address. The IP address for www.linuxformat.co.uk, for example, is 212.113.202.71. Machines on a small home network usually use IP addresses in the range of 192.168.0.1 to 192.168.0.256, though as you're administering it you can pretty much choose whatever you like!

MAC (MEDIA ACCESS CONTROL) ADDRESS Every network card connected to the network has a MAC address that is fixed to the device. Some firewalls work by limiting network access by MAC address – this is especially useful in wireless networks.

FIREWALL a hardware or software solution which monitors and prevents attempts to access the network from external sources such as the Internet. Linux has firewall features built in and the popular distributions have graphical tools for managing them. Look into the Network or Security sections. Most have basic options designed for setting up single machines connected to the Internet, network workstations or network servers. You shouldn't really have to venture beyond these settings unless you have very specific needs. If you find setting up a firewall a little too daunting, have a look at *Guarddog* (www.simonzone.com/software/guarddog) a useful configuration tool which is simple enough for basic use with power to spare if it is needed.



Not got a firewall? Check out *Guarddog*.

IP FORWARDING With this set up, a single connection to the Internet can be used by multiple machines. All traffic with the Internet will look, to all intents and purposes, to be enacted with a single IP address, even if it comes from a PC behind the gateway.

GATEWAY An intermediate between a computer and the Internet, used for the purposes of sharing a network connection. On a basic peer-to-peer network, one machine connected to the Internet can act as a router for other machines.

WIFI The consumer branding of 802.11a/b/g wireless networking standard.

NEXT MONTH

Amazing as it may seem, not every computer comes with Linux installed! Next month we will be looking at how other operating systems can be used to access and share data with Linux, and how Linux can be coaxed into running applications designed for other formats.

Apologies for the KDE bias in this tutorial. Everything covered here should be equally applicable to GNOME. For reasons of space, we have to concentrate on one or the other and, as a personal preference, the author tends to favour KDE – at least at the moment!



FLEX & BISON

Compiler writing

There's lots to learn when writing your own compiler, but, every so often, you need to go back over the code you've written to make sure it's in good shape. **Paul Hudson** does just that...



Example code from this tutorial can be found on your CD or DVD

A proverb states that a fool can throw a stone into the water that ten wise men cannot recover, and, although it was clearly around way before SKYLang, the same rules still apply – if we proceed on the course of adding more and more features, the chances are that the compiler will go beyond the point of no repair, leaving us in trouble if (when?) something breaks.

At this point in time, we have a fully working compiler and interpreter, working with our own brand of intermediate code. We also have loops and conditional statements working, which isn't bad considering that this is only part seven of the series. So far, though, we have mostly been adding new features to the

compiler and new syntax to the language, as this generally makes much more interesting reading. Having said that, as you will have seen if you've been trying out the code yourself – and you are, right? – you will have noticed that more and more parts of the compiler are not working as we expect them to.

These are generally small things; things that would go unnoticed if you only entered the SKYLang scripts presented in the tutorials. However, if an outsider were to try to write a SKYLang script, the chances are it would make the compiler segfault – not the reaction we want! So, the goal of this issue's tutorial is to try and close some of these holes in the compiler. Note that I say "some" – the chances are that there are dozens

more bugs that I either have no time to fix at this point, or aren't major enough to be fixed, or bugs that no one has spotted yet. We'll get to those in due course, but right now there are four particular problems with the language that I want to address.

Loop decrementing

The simplest place to begin is back at the loop code, as it was covered back in *LXF45*, and I also made the somewhat lame promise that I would pick up on the bugs in the loop code and fix them in *LXF46* – sorry! Anyway, better late than never, so the first bug we're going to fix in SKYLang is the fact that we can't have decrementing loops – our loops only go upwards.

To fix this problem, we simply need to duplicate the code for incrementing loops. This is a short-term fix, as code unification is something else we'll be looking at later in this issue. There are two pieces of code you need to duplicate: the first is the parser rule, and the second is the op code case. So, copy the following line of code, as well as the code associated with its action:

```
T_FOR T_VARIABLE T_ASSIGNEQUALS expression T_TO
expression T_INCREMENT expression T_OBRACE {
```

Now paste that in, changing **T_INCREMENT** to **T_DECREMENT**.

You also need to change this line within the duplicate:

```
SKYVar *result = sky_do_binop(SOP_ADD,
OpArray.opcodes[opnum]->op1, $8);
```

Note that it should read **SOP_SUB** in the copy, as we're decreasing the value of our iterator by subtracting from it with each loop. Once you've done that, we need to add some code to handle the comparison done by the decrementing loop. That is, a decrementing loop should terminate when the iterator is less than the minimum value, so our jump condition should be "greater than or equal to". So, change the parser action so that **new_op->opcode = SOP_JMPLTE;** becomes

```
new_op->opcode = SOP_JMPGTE;
```

So, we need to add the op code handler for **SOP_JMPGTE**. This is also a copy of existing code, **SOP_JMPLTE**, with the check **<=** changed for **>=**, and is therefore another easy change – check the source code on your coverdisc if you're not sure.

Non-executing loops

Back in *LXF45* I posed the following 'problem script':

```
$main = "Main loop";
$inner = "Inner loop";
for $i = 1 to 10 increment 2 {
  echo $main;
  for $j = 25 to 20 increment 1 {
    echo $inner;
  }
}
```

My question back then was, what was wrong with the script? Go ahead and compile your SKYLang build right now and try the script out, if you haven't done so already. The problem lies in the inner loop: if you do something once counting upwards for every number between 25 and 20, you should in fact never perform the action. However, because the nature of our loop-generation code is such that it performs the iteration check *after* the code has been executed once, the SKYLang interpreter will always execute a loop action, even if it is impossible.

The solution here is to run the opposite check before the loop actually takes place, that is, if the loop ends with a **SOP_JMPLTE**, you'd need to use a **SOP_JMPGT** before the first iteration. This

needs to be done after the variable is assigned so that it can be compared with a basic operation, but it needs to be done before the op number is pushed onto the loop stack, as we don't want to continually execute this statement.

Here's how that looks:

```
OpArray.opcodes.push_back(new_op);
new_op->opcode = SOP_JMPGT;
new_op->op1 = newvar;
new_op->op2 = $6;
new_op->result = new SKYVar;
LoopStack.push(OpArray.opcodes.size() - 1);
```

Note that the red lines of code are old – you need to insert the light blue code between the two lines shown here in red code in your copy. The new code is simply a jump if the loop variable is greater than the maximum value it can be – this is for incrementing in the loop. You need the opposite for decrementing the loop, and you'll also need to add **SOP_JMPGT** and **SOP_JMPLT** into *skylang.h*, as well as into the debug function **optostr()** if you have it in your code. Finally, you need to add cases for **SOP_JMPGT** and **SOP_JMPLT** into the **execute()** function, but these are easy.

So, with the new loop code in there, the problem script looks like this once converted to intermediate code (this will be in the file *skydump* if you have the debug function in your code):

```
Op 0: SOP_ASSIGN $main Main loop 0
Op 1: SOP_ASSIGN $inner Inner loop 0
Op 2: SOP_ASSIGN $i 1 0
Op 3: SOP_JMPGT $i 10 13
Op 4: SOP_ECHO $main 0 0
Op 5: SOP_ASSIGN $j 25 0
Op 6: SOP_JMPGT $j 20 10
Op 7: SOP_ECHO $inner 0 0
Op 8: SOP_ADD $j 1 0
Op 9: SOP_ASSIGN $j 0 0
Op 10: SOP_JMPLTE $j 20 6
Op 11: SOP_ADD $i 2 0
Op 12: SOP_ASSIGN $i 0 0
Op 13: SOP_JMPLTE $i 10 3
```

As you can see, operation **6** has our **SOP_JMPGT** jumping to operation **10** if **\$j** is greater than 20, which means it will jump to the loop iteration operation. Remember that our **execute()** function always pushes the operations forward one place, which means this will jump to operation **11**, continuing the **\$i** loop – perfect!

Constant strings

Take a look at the problem script again, this time confident that the inner loop won't be executed:

```
$main = "Main loop";
$inner = "Inner loop";
for $i = 1 to 10 increment 2 {
  echo $main;
  for $j = 25 to 20 increment 1 {
    echo $inner;
  }
}
```

The focus this time is the first two lines: they are set to strings, just to be echoed out again inside the script. Wouldn't the following script be more sensible?

```
for $i = 1 to 10 increment 2 {
  echo "Main loop";
  for $j = 25 to 20 increment 1 {
```



```

<< echo "Inner loop";
}
}

```

Go ahead and run that script through SKYLang. If you've been following so far, your output should simply be "Segmentation fault" – catastrophic failure! The reason for this is because our `execute()` case statement for `T_ECHO` reads from the `UserVars` map, without first checking whether the thing to echo is indeed a variable or whether it's a static value, as in the modified script above – that is, even using `echo 5`; will segfault, because SKYLang automatically treats all echoed data as variables. This is quite a basic flaw, and luckily also easy to fix.

The problem lies in this block of code:

```

case SOP_ECHO:
switch (UserVars[op->op1->charval]->type) {
case svInt:
printf("%d\n", UserVars[op->op1->charval]->intval);
break;

case svString:
printf("%s\n", UserVars[op->op1->charval]->charval);
break;

case svFloat:
printf("%.2f\n", UserVars[op->op1->charval]->floatval);
break;
}

```

Note that the first operand to echo, the value to print out, is automatically assumed to be a key in the `UserVars` map. As `UserVars` stores pointers to SKYVars, the map will return a null pointer if the key doesn't exist. As we attempt to switch on type without first checking whether we have a usable return value, this will naturally go wrong if we use a non-existent key. So, the solution here is to check whether we get a null pointer back from `UserVars` before trying to use it. Here's how that looks:

```

case SOP_ECHO:
if ((UserVars[op->op1->charval]) != NULL) {
switch (UserVars[op->op1->charval]->type) {
case svInt:
printf("%d\n", UserVars[op->op1->charval]->intval);
break;
case svString:
printf("%s\n", UserVars[op->op1->charval]->charval);
break;
case svFloat:
printf("%.2f\n", UserVars[op->op1->charval]->floatval);
break;
}
} else {
switch (op->op1->type) {
case svInt:
printf("%d\n", op->op1->intval);
break;
case svString:
printf("%s\n", op->op1->charval);
break;
case svFloat:
printf("%.2f\n", op->op1->floatval);
break;
}
}
}

```



```
break;
```

That's the entirety of the change, so it's nothing too tricky. As you can see, the code now checks to see whether it gets a null pointer back from `UserVars`, and, if it does, it outputs the value of the variable directly, without trying to grab any values from `UserVars`. The actual code to output variable contents is almost the same as if we have a variable, with the exception that we work with the value directly, as opposed to using its charval value to index into `UserVars`.

The modified script, echoing out strings directly rather than setting them as variables first, should now work perfectly. However, that's just the tip of the iceberg, because as a result of the code change above, more complicated statements can be written, like this:

```

$foo = 5;
echo $foo + 10;

```

Two birds with one stone, if you like.

Operator unification

Ideally we'd like our if statements to be able to handle all sorts of conditionals, but it's a little overkill to have to replicate the parser actions for each type of equation – after all, the only difference between an if statement checking whether one number is higher than another number is the operator in the middle. As such, the ideal solution is to be able to create a generic parser action that handles various different operators, and this takes a little thinking to accomplish.

The annoying part is to set up the execution cases for each type of operation, so you should get that done first. You should already have `SOP_JMPLTE` working for all types of variables; the next step is to duplicate that case for the other types of operators, changing the code as appropriate. The changes are obvious: for `SOP_JMPGTE`, change `<=` to `>=`; for `SOP_JMPLT`, change `<=` to `<`, etc. This has all been done for you in the source code on the coverdisc if you want to skip the hassle.

The more interesting part is where we handle the different operators in the parser. This is accomplished with a whole new non-terminal, called `comp_operator`:

```

comp_operator: T_ISEQUALS {
    $$ = SOP_JMPEQ;
}
| T_LESSTHANEQUALS {
    $$ = SOP_JMPLTE;
}
| T_GREATERTHAN {
    $$ = SOP_JMPGT;
}
| T_GREATERTHANEQUALS {
    $$ = SOP_JMPGTE;
}
| T_LESSTHAN {
    $$ = SOP_JMPLT;
};

```

All that does is convert the tokens into the appropriate jump statement. Note that it returns a `SKYOpType`. To make that work properly, you need to amend the `Bison` union to this:

```

%union {
    SKYVar* var;
    int intval;
    double floatval;
}

```

```
char* charval;
SKYOpType optype;
}
```

Note that it can now contain **optypes**, which is where we'll be storing our converted tokens. Just below the line **%type <var>**, add these lines:

```
%type <optype> comp_operator
%token <intval> T_LESSTHANEQUALS
T_GREATERTHANEQUALS T_LESSTHAN T_GREATERTHAN
T_ISEQUALS
```

Note that the return value of our **comp_operator** non-terminal is declared as a **SKYOpType**, as planned. The **intval** for our comparison operators is just nominal. Here's the important part: now that we have the ability to convert comparison tokens into jump operations, we now just need to amend the master if statement action to take advantage of the new non-terminal, like this:

```
T_IF T_OBRACK expression comp_operator expression
T_CBRACK T_OBRACE {
SKYOp *new_op = new SKYOp;
OpArray.opcodes.push_back(new_op);

new_op->opcode = $4;
new_op->op1 = $3;
```

As **comp_operator** returns a **SKYOpType** directly, we can just provide it directly to the opcode as **\$4**, the fourth part of the action match rule – perfect!

Updates and concatenation

The operation unification above is nice enough in itself, but there's still more to be done. If you recall from last issue, a lot of the execution code still tries to call **op->op1** and **op->op2** directly, which circumvents the new solution we implemented that automatically retrieves variables as necessary. So, the first change to make here is to replace all instances of **op->op1** and **op->op2** in the **switch(op->opcode)** case list – note that you should make sure you only replace inside that selection, otherwise you will break a lot of code.

Once that's done, you need to make one more tweak right now. You see, the code we implemented earlier in this article to allow string constants to be echoed now needs to be changed again, because the **SOP_ECHO** case tries to look up the operands to see if they are variables that require replacing – an action we do by default now with **tmpop1** and **tmpop2**. As such, you can amend **SOP_ECHO** to the following very easy piece of code:

```
case SOP_ECHO:
switch (tmpop1->type) {
case svInt:
printf("%d\n", tmpop1->intval);
break;
case svString:
printf("%s\n", tmpop1->charval);
break;
case svFloat:
printf("%.2f\n", tmpop1->floatval);
break;
}
break;
```

Once that's done, you should now be able to make the compiler, this time with the new code. The other change I wanted to make was to add variable concatenation – quite a big hole in our current operator list. Concatenation can be quite a lengthy topic to teach, so I hope you'll forgive my forced brevity!

To get started, you need to add the following lines to the lexer in **skylang.l**, just below **T_MODULO**:

```
“:” {
return T_CONCAT;
}
```

Open up **skylang.h**, and add **SOP_CONCAT** into the list of **SKYOpTypes**. The next step is to add the **execute()** case for **SOP_CONCAT**. Here it is:

```
case SOP_CONCAT:
if (tmpop1->type != svString) { tmpop1 = sky_tostr(tmpop1); }
if (tmpop2->type != svString) { tmpop2 = sky_tostr(tmpop2); }

op->result->charval = new char[strlen(tmpop1->charval) +
strlen(tmpop2->charval)];
strcpy(op->result->charval, tmpop1->charval);
strcat(op->result->charval, tmpop2->charval);
op->result->type = svString;

break;
```

As you can see it's a great deal easier than the other operations, simply because we only concatenate strings – everything that isn't a string gets converted to a string using **sky_tostr()**, a function we'll be looking at momentarily. Using **strcpy()** and **strcat()**, the concatenated strings are placed into the result of the operation, and sent back.

The **sky_tostr()** function is also very simple, thanks to C's **sprintf()** function. Here's the code:

```
SKYVar* sky_tostr(SKYVar *var) {
SKYVar* working = new SKYVar;
working->charval = new char[255];

if (var->type == svInt) {
sprintf(working->charval, "%d", var->intval);
return working;
}

if (var->type == svFloat) {
sprintf(working->charval, "%.2f", var->floatval);
return working;
}
}
```

As you can see, it simply runs a basic conversion of types, and returns the new value in the working variable. This is almost the same as the **sky_strtoint()** function, so it shouldn't need any explanation.

Conclusion

While our compiler is still a long way from perfect, we have at least managed to nail down some of the more obvious bugs and glitches, as well as also tidy up the code. From here we are on stable enough ground to be able to proceed with new features (hurrah!), namely function calls. This signals our return to the big features push, and you should put your thinking caps back on because it's *not* going to be easy. But, then, it never is! [LXF](#)



NEXT MONTH

We'll actually be taking a short break from this series next month, but it will return in issue 49! How do functions fit into the SKYLang language? What built-in functions should we make available to the legions of programmers waiting to be introduced to SKYLang? It's all revealed in a special sixteen-page Compiler Writing tutorial in the next issue! Well, give or take twelve pages probably...

ILLUSTRATION BY MARCUS STADE - www.blender3d.org

ADVANCED EFFECTS

The Blender Game Engine

PART 8 Jono Bacon monkeys around with the interactive game construction facilities.



Throughout this *Blender* series, we have been using the software primarily for modelling non-interactive scenes and animations. Although this is the main area in which *Blender* is generally used, there is another side to *Blender's* personality; the Game Engine. Wide-eyed readers may now be quite excited at the idea of making the next *Quake* or *Resident Evil* clone, but what does the Game Engine actually allow us to do, and do we need a degree in mathematics and physics to use it?

Luckily for fans of our beloved modeller, the Game Engine is not only simple to use, but is incredibly powerful. The basis of the Engine is to allow us to create interactive 3D graphics that are generated in real-time as opposed to being pre-rendered. In English, this means that the Game Engine has to create all of the graphics on-the-fly as opposed to creating them first and saving them to a file as we do when we render. In reality, this means lower quality graphics if you want your game to run on anything other than a supercomputer. Although the graphics are cut down, they are by no means shabby, and you can get results

comparable to that of most games consoles and PC games; the limitation is merely your hardware.

Before we get a too excited, I must stress one important requirement for using the *Blender* Game Engine. Since *Blender* was open-sourced, the Game Engine was removed due to some reliance on non-Open Source software. Although the licensing issues have been resolved, the last version of *Blender* to have the game kit was version 2.25 (the most recent version being 2.28 at the time of writing). You will need the 2.25 version to make use of this tutorial.

Getting started

To get us started with the Engine, we will first create a simple scene. Create a new project and instead of removing the default plane as we normally do, resize it so it is much bigger. When the plane is a reasonable size, add a shape to the scene such as an icosphere. There is a special shape included with the 2.25 version of *Blender* called *Monkey* which creates a monkey's face

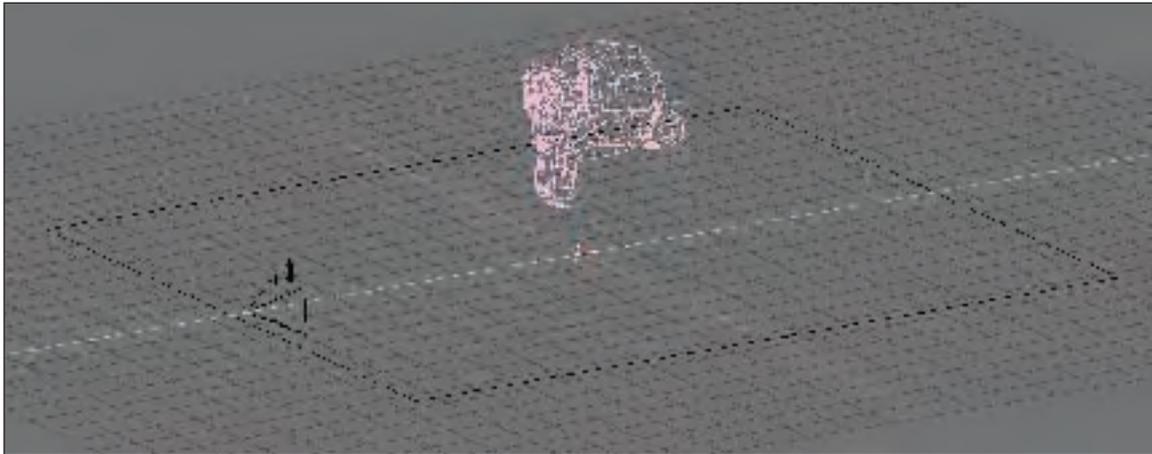


Fig1 Our monkey hovering above the ground.

model. I have added the monkey shape, and hovered it over the flat plane. Using both mouse buttons together, rotate the view so that you can see the two shapes in a more three-dimensional way. Your screen should now look similar to **Fig1**.

With our impressively simple scene set, we can now begin to set the monkey to be a character in the game. To do this we will make use of the Realtime buttons that can be accessed with the **F8** key. The concept is that we first assign an object to be a character (referred to as an Actor in *Blender*), and then we will determine what physical constraints the actor has associated with it (such as gravity, weight, speed *etc*).

In the real-time buttons, at the top left is an Actor button. Select the monkey with the right mouse button, and then press this Actor button. When you click the button, two more will appear; Ghost and Dynamic. These two buttons allow you to specify what type of character our monkey will be. A Ghost refers to an entity that does not collide with other entities and has no physical constraints, and a Dynamic character has a physical form that is affected by the physics of the scene. Physics plays an important part in most games (even simple physics that dictate that the scene has earth-like gravity), so we will set our character to be a Dynamic actor. Click on Dynamic and a number of other buttons appear.

The buttons that are shown allow you to create your basic physics for the scene. If you view the scene now (press the **P** key), you will see that the monkey falls from the sky. Although the monkey falls and bounces on the plane satisfactorily, it could look better with a bit of tweaking. We can adjust this behaviour by using the Damp button. If you increase its value to 0.800, you will see that the object seems to have considerable weight now, and the bouncing becomes more realistic. An important setting that you should configure in these buttons is the Size button. This button sets how much of the shape can collide with other objects and you will see a dotted circle around the shape. Increase the size of the dotted circle by adjusting the button, and ensure that it is the size of the monkey.

To test our collision, we can add another object to our scene that our monkey could collide with. This can be done by adding a shape such as plane underneath the monkey, so it will bounce off it when falling down. I have added a plane and rotated it so it forms a ramp, as shown in **Fig2**.

You will see that when you play the game, the monkey will bounce down off the plane and land on the ground. This shows how different objects can interact with each other, even at this basic level.

Controlling your characters

It is very unlikely that you could find a game that has no interaction from a user – interaction is the key to fun in a game. Interaction in the *Blender* Game Engine is handled by a number of different facilities within the real-time buttons (**F8**). Although setting an interaction can seem a little complex at first, it is not difficult to get results with some practice.

To get us going with some interactions, remove the obstacle that you created and put the monkey on the default plane so it is resting on it. To add interaction to the scene we need to first select the object that can be controlled; in our case, this is the monkey, so select it with the right mouse button. With the object selected, now go to the realtime buttons (**F8**) and you will see three blocks of controls to the right of the buttons. These are labelled Sensors, Controllers and Actuators. In these sets of buttons we will configure our interaction.

First, click on the Add button next to the Sensors label, and a lot more buttons will appear. This block of buttons lets us configure which sensor can trigger something. In our case, we want our sensor to be a key on the keyboard (such as the Up Arrow key, for instance), and we would like to move the object forward. Each of the Sensor, Controller and Actuator boxes has an X icon to delete the current entry, and an arrow to maximise/minimise the box.

In the sensor box, we have a drop-down box with Always written in it. This drop down box lets you select the type of sensor that you will be using. There are many different types of sensors for different interactions from the user, and different events within the Game Engine. We would like to control the monkey using the keyboard, and you can select Keyboard as the sensor from this box. You will then see that some more entries have been added to the box. The most notable of these is the Key part. If you click in the area next to where it says Key, you will be instructed to

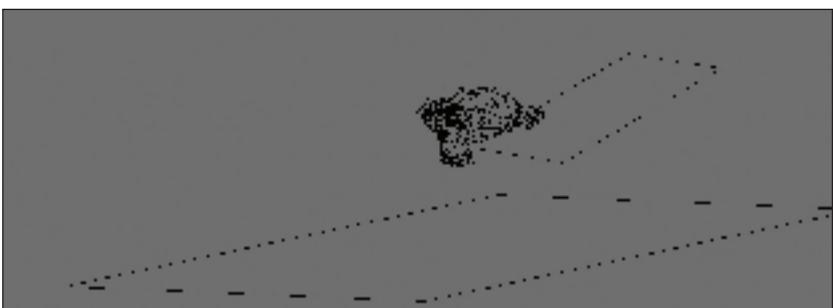


Fig2 Putting objects in the way of a monkey and a plane.

CREATING A DRIVING GAME

Scratchbuild your own *Gran Turismo!*

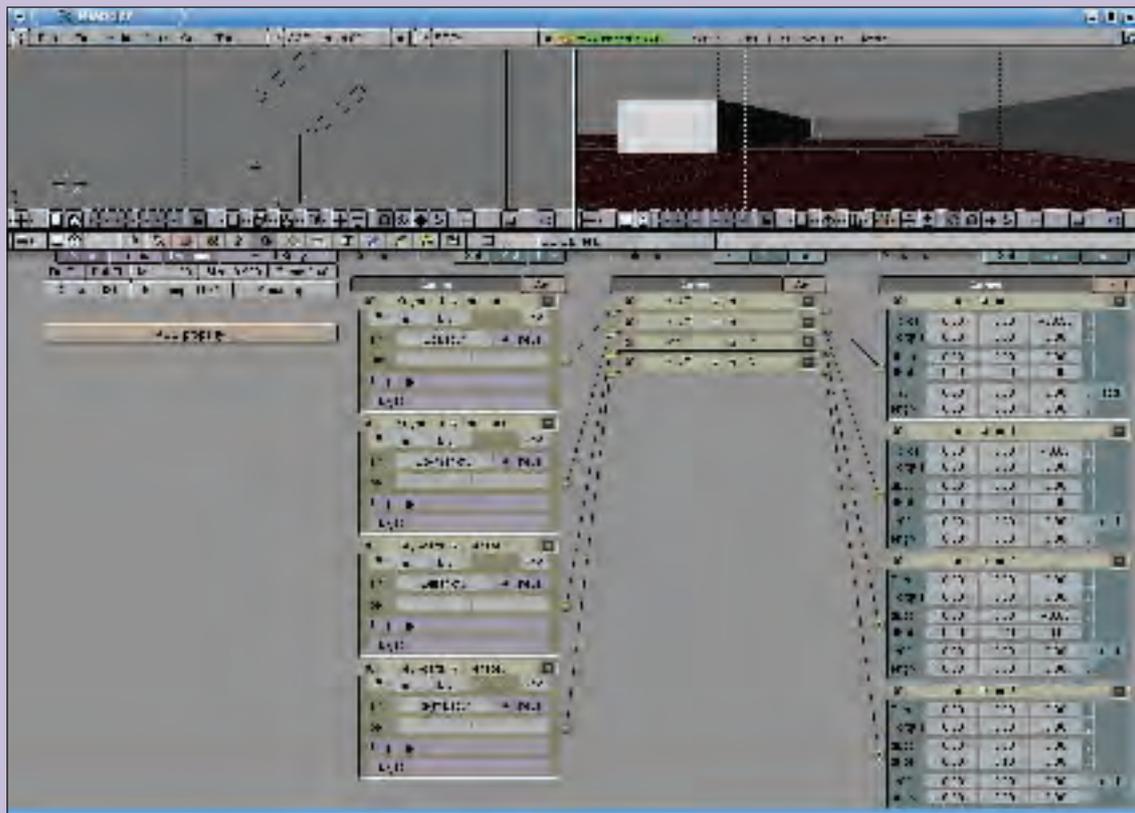


Fig3 The real-time settings for the game.

With the concepts that we have discussed in this issue, we can create a simple first person perspective driving game. First, create a new scene and increase the size of the default plane so it is quite large. Next, split the views so that you have a camera and top view, and add a number of cubes that will form the outside barriers of your simple road. Next, position the camera at the start of the

mocked up road and select the camera with the right mouse button. Using the real-time buttons, set the camera as an actor and create the Sensor/Controller/Actuator relations for each arrow key.

You will need to set the up arrow to go up an axis, the down arrow to go down an axis, the left arrow to rotate left slightly, and the right arrow to rotate right slightly. When each

of the arrow keys has been configured you can now hover the mouse over the camera view and enter the game mode by pressing the **P** key. You will now be able to drive around the course, and you will rebound off the walls. If you replace this basic functionality with better-looking scenes and a mock-up of a car in front of the camera, you have your very own driving game.

◀ press a key that will act as a sensor. In our case, press the Up Arrow key and you will see it added to the box.

The next section is the Controller part, and you can click on the Add button to expand the box and create a controller entry. This box basically allows you to combine multiple boxes together to form complex sensor/actuator relationships. We do not need to do any complex combinations at the moment, so you can leave it as an AND setting. To link the Sensor box with the Controller box, we can click on the yellow ball to the right of the Sensor box and drag it over to the yellow hollow ball to the left of the Controller

box. You will see a line appear to link the two boxes together. This relationship can be removed by clicking on the line and pressing the **X** key.

The next and final section is the all-important Actuator. When you click Add to create an actuator, this section will do the hard work of adjusting your model so that it will change dependent on your key-press. The box has a number of sections that refer to different effects on the model. As an example, the Force line will adjust how the object is pushed. The dRot line is also important for rotating the object. Each of these lines has three numbers

WHEN WILL THE GAME ENGINE BE BACK?

Set for inclusion in future *Blender* iteration

As was mentioned at the start of this article, the Game Engine was removed from the main *Blender* program when the software was open-sourced; the reason for the removal being the collision detection software that *Blender* uses in the Engine.

Although the Engine was removed to stop licensing problems, the developers got on with developing a replacement library to bring the Game Engine to the new Open Source *Blender*. Shortly after work begun on the replacement, there was news that the

original library used by *Blender* had been open-sourced as well, and the Game Engine looks set for a return soon with even more features. The specific release with the new Game Engine has not been announced yet – *LXF* will let you know as soon as it happens!

NOT JUST FOR GAMES

Solid 3D functionality across the board

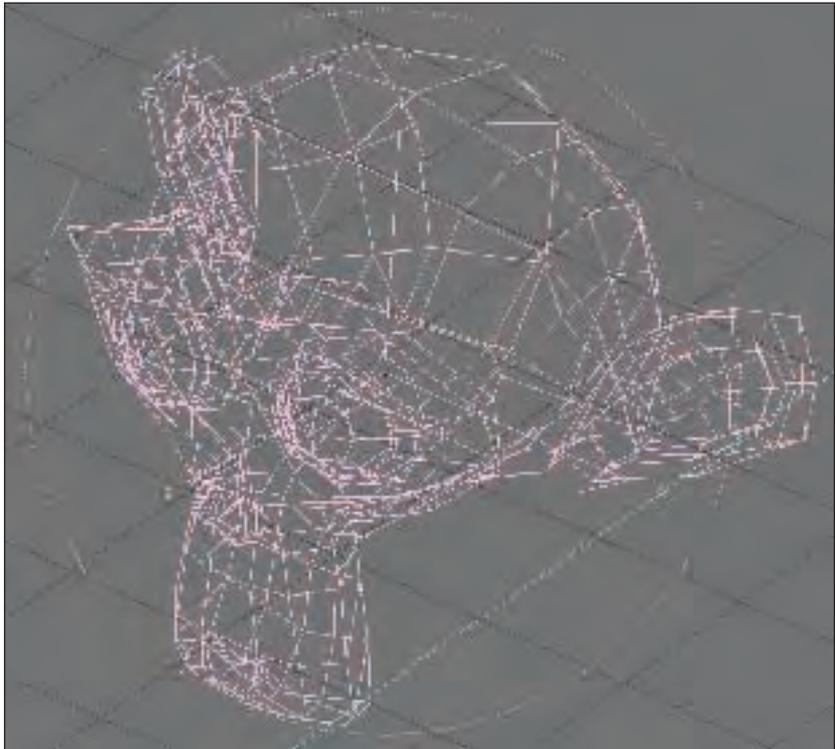
One of the most deceiving aspects about the *Blender* Game Engine is the fact that it is called the "game engine". Admittedly, many people make use of the Engine to create interactive games and puzzles, but there are many other uses for the Engine in more conventional settings. Many of these applications of the Game Engine can be used in the workplace, and not only does *Blender* provide a solid 3D platform, but an interactive graphics facility is present too.

The typical use for the Game Engine is with interactive walkthroughs. As an example, if you were an architect, you could make use of the Engine to literally walk into the building and walk around it. This would vastly improve on the typical linear animated walkthroughs

common with building design. Another use would be for walking around a product such as a car. The interactive nature of the Game Engine could mean that all angles of the car can be seen at any time; just like walking around a showroom.

Blender can also be used to create virtual worlds and site plans. With enough time, you could model an entire town in the Game Engine and move around the town exploring the different streets. This can be used for creating maps, guides, layout plans and other uses.

Although a cliché, your imagination really is your limitation, and anything that you can visualise and interact with in the real world, could be reasonably implemented in the Game Engine.



that refer to the X, Y, Z axis of the object. You can see the axis for your object by going to the Edit buttons (**F9**) and clicking on the Axis button on the left hand side. **Fig4** shows the axis being shown for my object.

With my object, the direction that the monkey is facing is the Z Axis. To move the monkey forward, I can now increase the value of the third column in the Force row (the Z axis column). If the monkey was to go back (when it is assigned to the Down Arrow), the number in the Z column need to be a negative number; the number is relative to the objects current position. To finish off, connect the Controller and Actuator boxes together with the line (in the same way we connected the Sensor and Controller boxes earlier). When you now enter the game mode by pressing the **P** key, the Up Arrow will move the monkey forward.

It is important to remember that the axis of the object is always relative to the position of the object itself. In other words, if the object is facing a different way, the Z axis will face the way the object is facing; the axis moves with the object. In this sense, once you have set a key for the monkey to go forward, it will always move the monkey forward, no matter which direction the monkey is facing towards. You can now set the other keys.

Conclusion

In this issue we have taken a look at the basic elements of using the *Blender* Game Engine. The Game Engine is not only a simple piece of software, but it is an extensible component in *Blender*, and there are many aspects to it that we simply do not have the space to cover. If you are interested in taking your knowledge further, it is recommended that you first have a look at the many existing games available, and consider buying the Official Gamekit book that is sold by the *Blender* project (your purchase money goes to help further development of *Blender*).

One major feature in *Blender* is its support of the Python language. This language has been added to allow you to script events and facilities in *Blender*, and the language really comes

into its own in the Game Engine. *Blender* offers the best of both worlds; you can create simple games with no programming using the *Blender* graphical tools, but can create advanced games using the Python support in the environment. One other aspect of the *Blender* Game Engine is the ability to play the games without *Blender*. This is possible using the now free *Blender Publisher* (it used to be a commercial product), and will allow you to create stand-alone game applications.

In this *Blender* series we have covered an extensive amount of ground in a relatively small number of pages. From our humble beginnings eight months ago when we started by adding a shape to a scene, we have covered different shapes, lighting scenes, materials, textures, advanced modelling, modelling with curves, key frame/bone animation, particle systems and now the Game Engine. Even though we have spent time introducing all of these concepts, there are still many things to learn in *Blender*.

Blender is an incredibly feature-packed application and it has a very bright future because it is constantly under development. A lot of developers are now involved with producing brand-new functionality and features for the fully Open Source *Blender*, and the with each addition, *Blender* becomes more able to hold its own in a market that is still dominated by proprietary applications. Keep on practising, tuning and exploring these new features, and keep on blending! **LXF**

Fig4 The axis connected to my monkey object.

NEXT MONTH

Over the last eight months, we have covered many basic ways of making scenes. In the next issue of *LXF* we will examine how some of these techniques can be used in order to add extra realism to your scenes in *Blender*.

This series is drawing to a close quite soon, so if you've got any last-minute suggestions for *Blender* methods or questions that you would like to see covered in the next couple of issues, please email them to linuxformat@futurenet.co.uk without delay!

IMAGE EDITING

Custom patterns and brushes

Michael J Hammel steps back just a bit to look at some basic of the basic building blocks of successful art projects: creating your own patterns and brushes.



The articles offered in this *GIMP* tutorial series have tried to look at some real-world issues as opposed to the more mundane user interface questions, but there are probably a few folks still wondering about some more basic operations. If the *GIMP* is the construction set for Linux graphic artists, then surely patterns and brushes are the basic building blocks. Few projects can completely evade the use of these two fundamental features. Yet as useful as they are, each can be made even more meaningful to the individual artist. All that is required is a little creativity and basic knowledge of the *GIMP* user directory structure.

Pattern files are nothing more than ordinary image files saved in *The GIMP's* native pattern format file, known as PAT files. These files tend to be rather large, larger than TIFF files and significantly larger than compressed JPEG files. If you want to distribute large numbers of pattern files you could save a little room on the CD or your Web site by distributing the TIFF versions instead and letting your users convert them individually

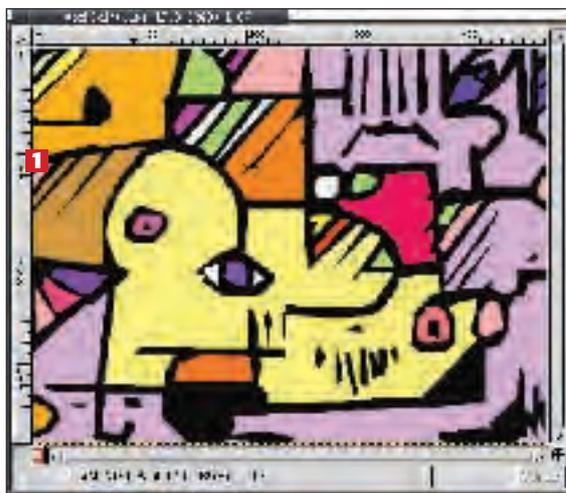
to pattern files. TIFF and JPEG are also commonly supported by image viewers on many platforms, making the patterns easy to browse before installing. The PAT file format is supported only by *The GIMP* at this time.

Brushes come in multiple formats, all of which are native only to *The GIMP*. Brush files from *Adobe Photoshop* are not compatible at present, but creating your own versions of those brushes is easy. Just paint a single brush stroke on an image in *Photoshop*, save as an uncompressed TIFF, load it into *GIMP* and save it as a brush. This methodology allows you to create large collections of your own personal brushes, just as you might collect brushes of varying shapes and sizes for your oil and watercolour paintings.

This issue we'll look at creating both patterns and brushes and the simple steps required to save images in the proper format. We'll also look at Brush Pipes, a powerful yet seldom discussed type of brush that allows you to use multiple brush tips at the same time.

Making Seamless GIMP Patterns

The GIMP's set of stock patterns is a reasonable, but far from complete. Typical pattern users – such as those working on textures for 3D models – often find themselves creating and managing their own personal set of patterns. Creating a pattern is simple enough: there are no ground rules for what a pattern must be other than the file format used to save the file (*The GIMP* native PAT format).



Open clipart, colorized

1 We look through our collected set of normally useless CD archives to find a set of WMF formatted clipart. We can use these as the basis for our patterns. We want the pattern to fill the Canvas window and to be seamless when applied side-by-side with a copy of itself. This image shows clipart cropped to fill the Canvas window but not yet made seamless.

Make Seamless

2 There are two methods to make the image so it tiles seamlessly. The first is to offset the image by 1/2 (Image >Transforms>Offset). This option is shown in the upper left of this image. You can use the clone tool to fix the seams manually if necessary, which this image would surely need. The other method is to use the Filters>Map>Make Seamless filter. This will produce an effect that makes the pattern look slightly ghosted, but the resulting image can now be tiled seamlessly on web pages or in print images.





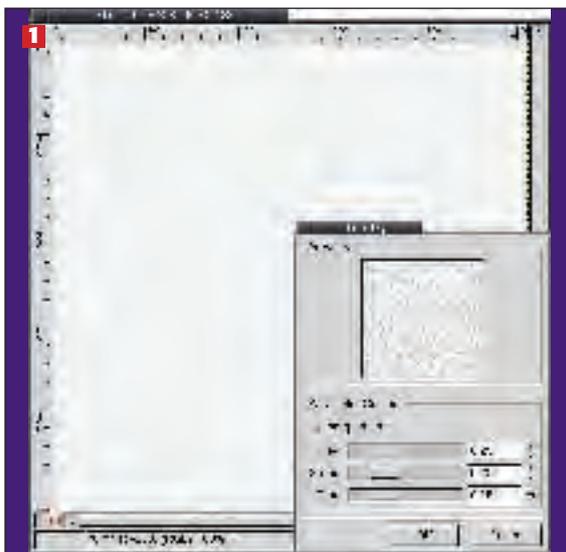
Save as pattern

3 Once you have your image made seamless, you are only left with saving it as a pattern. Choose the PAT format and give your pattern an easy-to-remember name. This name will be used in the Pattern Selection dialog. Patterns must be saved to your `.gimp-1.2/patterns` directory and the Pattern Selection dialog must be manually refreshed to see any newly installed patterns. Note that patterns can be any size. Larger patterns with high detail are better suited to 3D textures so that the repeating pattern isn't as noticeable.

Pattern Tips

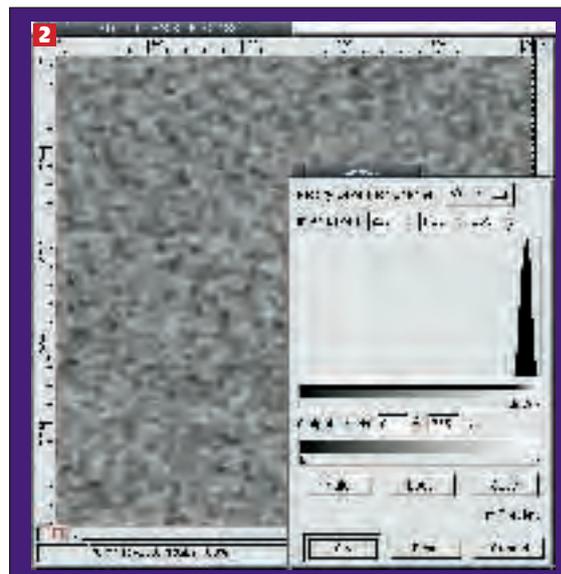
There are many uses for patterns. Background images for webpages are common, to the point of sometimes being annoying (though widespread use of CSS may change that a bit). An interesting alternative is to use patterns for displacement images. This allows patterns to act as 3D textures applied to another image. Examples of this include grayscale patterns that simulate water, skin and metal effects.

In the following steps, we'll create a lizard skin and show how cloning with freehand selections can be used to make the pattern seamless.



Noisified and blurred canvas window

1 We start with a white canvas. Using **Filters>Noise>Noisify**, we add 25% noise for each channel to the image. The noise is then blurred using **Filters>Blur>Gaussian Blur (RLE)** with a horizontal setting of 9 and a vertical setting of 7. The screenshot shows the noise filter dialog and the Canvas window after this noise has been applied and blurred (though with such low contrast that may be hard to see in print).



Level adjustments

2 The blurred image has a very short range of colour, as the Levels dialog shows. Adjust the black point to fit the left edge of the histogram or just use the Auto Levels option. This will bring out the details as shown here. Apply a **Bump Map (Filters>Map>Bump Map)** to this image using itself as the bump map image. Use 220 degrees of Azimuth, 62 degrees of elevation, and a depth of 16 for your bump map. This is what gives you the rough, 3D appearance of the skin.



Colorized with multiple layers

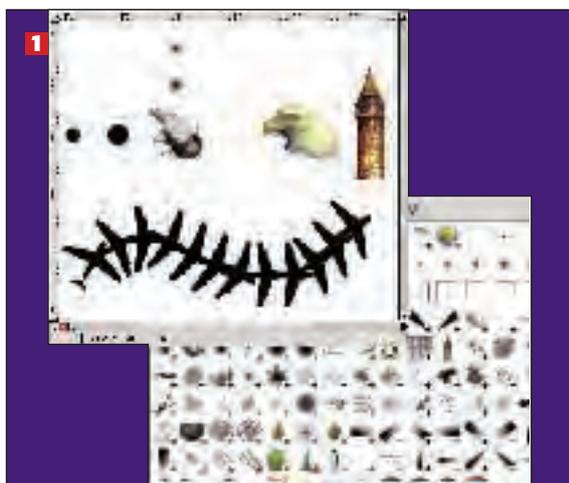
3 We add some green paint on a separate layer along with a golden gradient layer (the Layers dialog in the lower left of the screenshot) to add colour to the texture. Flatten this and then offset the image by 1/2 (**Image>Transforms>Offset**) and make freehand selections to cover seams (the upper right and upper left of the screenshot shown here). After we flatten the image and save it as a PAT file to our `.gimp-1.2/patterns` directory and refresh the Pattern Selection dialog, we're finally ready to apply the pattern to a larger canvas to check for seams and obviously repeating patterns (lower right of screenshot). >>

TUTORIAL GIMP

« GIMP brushes

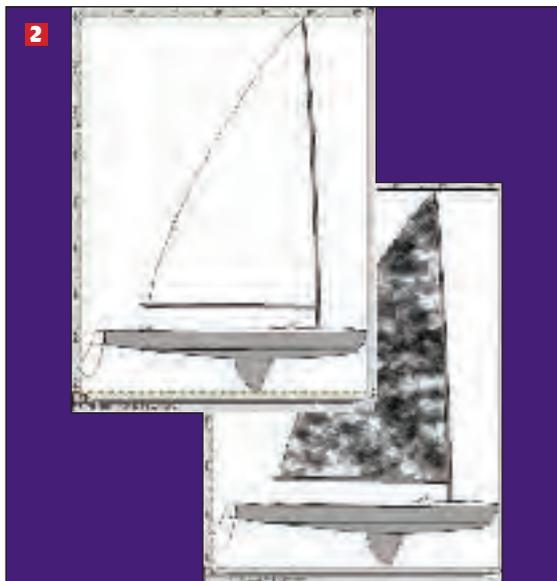
The *GIMP* is, above all, a sophisticated digital painting application that provides the artist multiple brush options, including soft and hard-edged paint and air brushes, pencils, and cloning and ink tools. Each of these tools, in turn, can be used with multiple types of brush tips. These brush tips (referred to generically as brushes by both official documentation and most users) provide variations on the pattern that is laid out while painting a single brush stroke.

Many users might not consider the value in the variety of brushes that can be used with each of these tools. Brushes can be configured for variations of spacing, opacity, blend modes, fade out and colour.



Multiple brush types in single canvas

1 There are three main types of brushes in *GIMP* – grayscale (the dots and pointing hand), colored (the eagle and castle) and Brush Pipes (the airplane images, which are actually just one brush pipe). The Brush Selection dialog (shown partially hidden behind the canvas window) shows all the available brushes. Brushes that are larger than they appear have plus (+) signs. Brush Pipes have small red triangles as well.



Loaded clipart and inverted grunge layer

2 Brushes allow us to easily add life to lifeless clipart. We've opened a clipart image of a sailboat, made a selection using the Fuzzy Select inside the Sail area, created a new layer and painted with a Grunge Brush. Grunge brushes are not in the stock *GIMP* distro, but you can use any brush you have that adds distinction to the sail. We inverted the grunge layer so that the painted area will be white. This makes it easy to apply a color blend mode layer above this to colorize the image.

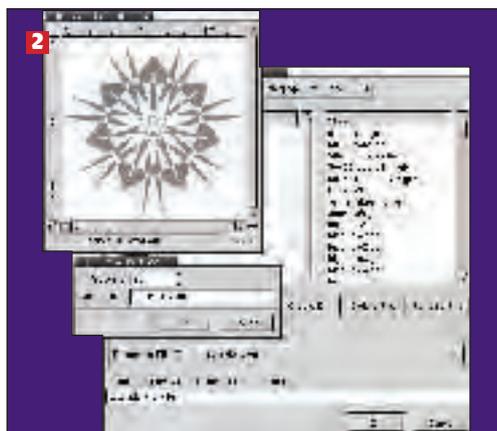


Final version

3 Leaving the selection intact makes it easier to paint only the sail area. In a new layer we've selected a Brush Pipe brush (made up of leaves) and set the Paintbrush 'Gradient' option. This allows us to paint in the new layer, and only inside the selection, with leaves that are coloured by the gradient. The sample image shows the layer we painted where the leaves are clearly visible. When this layer's blend mode is set to Color, the shape of the leaves is lost, leaving only the original grunge shape. Painting using content layers (the grunge layer in this example) and colour layers (the leaf layer) allows us to easily change the color OR the content without affecting the other.

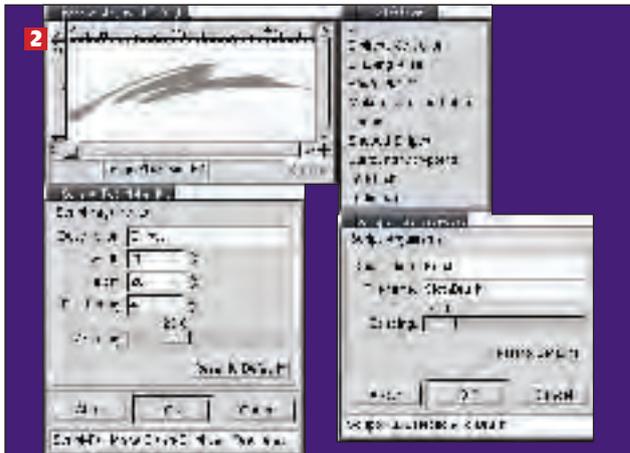
Creating and saving brushes

The standard set of brushes provided by *The GIMP* is good, but limited. Extended sets can be found on the Internet, such as the collection of over 150 brushes found in the *Graphics Muse Tools CD* (this is also where you'll find the Grunge Brushes mentioned in the previous tutorial). Creating your own grayscale or coloured brushes is practically a no-brainer. All that is required is a little creative thought and a few simple steps.



Clipart and save options for GBR

1 Any good clipart, properly sized, can be used as a new brush. Here a clipart in WMF format was opened, scaled to 35x35 and saved in the GBR format to the `.gimp-1.2/brushes` directory. The clipart can be coloured prior to saving as well, making it a coloured brush. The larger the image, the larger the brush. If you need a brush that is only a few pixels in width and height, try zooming in on it and edit the pixels manually. The image needs to be flattened and an alpha channel added before saving it in GBR format. The white background should be removed if possible otherwise you'll get a white square painted with each brush tip that is applied during a brush stroke.



Alternate menu items for saving brushes

2 Brushes can also be created quickly using a selection and the Script-FU>Selection>To Brush... option on the Canvas menu. Just make a rectangular selection around the shape you want to convert into a brush, choose the "To Brush..." menu option and fill in the dialog. The selection is saved to your .gimp-1.2/brushes directory and the Brush Selection dialog refreshed automatically. If you need a simple square or elliptical brush, try using the Xtns>Script FU >Make Brush menu option from the Toolbox.



Canvas with painted lines and updated brush dialog

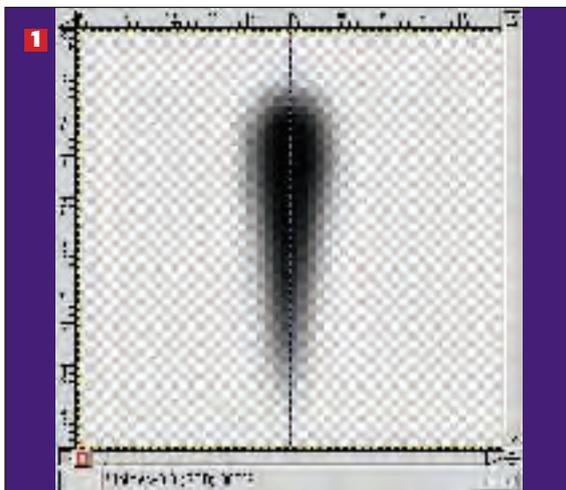
3 Changing the spacing for the brush will change the number of times the brush is applied during a brush stroke. In this image there are three paths to which have been stroked with the new coloured brush, each using a different spacing. The barely visible one on the left uses the default spacing for the brush (which was 20 in this case). The next uses a higher value of 55 and the last an even higher value of 170.

Brush (or Image) Pipes

A relatively new alternative to the static grayscale and coloured brushes are the Brush Pipes. Brush Pipes appear as normal brushes until you click and hold the left mouse button on them in the Brush Selection dialog. They are then expanded and present an animated appearance. However, these are not

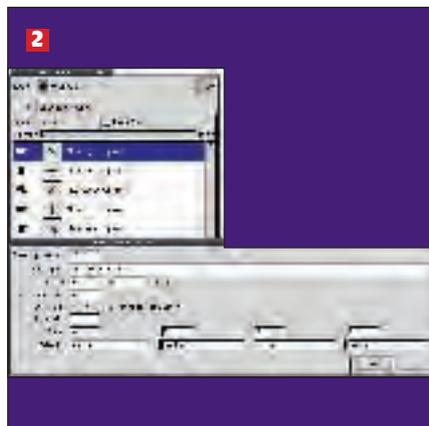
animated brushes. Instead, these are brushes with special characteristics that allow a different brush tip to be used based on a variety of options, such as the direction of the brush stroke or the pressure or tilt of a drawing tablet pen. Brush Pipes are created by building a multi-layered image, with each layer

holding one or more brush tips. The actual tip painted during a brush stroke depends on the options you set when you save your layers in the GIH format. In this tutorial we'll take a series of related brushes (the Quill brushes from the Graphics Muse Tools CD) and convert them into a single Brush Pipe.



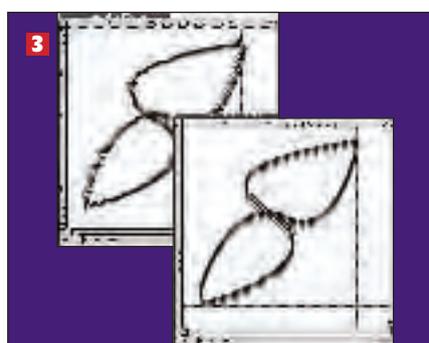
Canvas with grids and single brush stroke

1 The set of brushes we'll be using to create the Brush Pipe are all small enough to fit inside a 50x50 canvas. We start by creating a canvas of 50x50 pixels with eight transparent layers. Zooming in and using Ctrl-E to resize the window to fit allows us to better see what we're doing on such a small canvas. The Guides>Center Guides menu option was used to create a crosshair on the image. On the bottom layer we use the Paintbrush to paint a single brush tip.



Layers dialog -all layers

2 Repeat applying a single brush tip per layer, with each layer getting a different version of the 8 available Quill brushes. The order of the layers (ie the direction of each quill brush) matters only for certain options when saving the Brush Pipe. Then save the image with layers intact to .gimp-1.2/brushes directory using the '.gih' filename extension. In the 'Save as Brush Pipe' dialog, set the Ranks to 8 and Selection to Angular in the first column and leave other values at default. This produces a Brush Pipe that selects one of the Quill images based on the direction the mouse moves while painting.



Canvas with two curves, each w/ different brush pipe

3 If we create another Brush Pipe that sets the second Rank to 2 and second Selection field to Velocity we get different results when we paint. Here we see what two paths look like when stroked with the first Brush Pipe (Angular only) and the second Brush Pipe (Angular and Velocity). [LXF](#)

FASTER CODE

Practical PHP Programming

Sockets are good for all sorts of things beyond just connecting the iron into. Want to revolutionise your PHP scripts? **Paul Hudson** plugs you in...

Files are files, directories are files, and devices are files. Everything, and, conversely nothing (think /dev/null), is a file. So it should come as no surprise to you that sockets are also files, and the logical extension of that fact is that you can manipulate files and sockets in precisely the same way.

But first, let me explain what a socket actually is. Perhaps the best way to think of a socket is like a connector, and might be between a program and a port, or perhaps two programs – data goes in one end, and comes out the other end. There's a lot more to sockets beyond that, but it's not important on the whole – at least not for this installment.

File-like sockets

We covered files only very briefly, and even that was back in *LXF31*, so if functions like **fopen()** and **fread()** look alien to you, you'll need to dig out *LXF31* from the bottom of your cupboard. As seen back then in the second installment on Practical PHP, **fopen()**, **fread()**, **fwrite()** and **fclose()** are used in combination to

make file handling straightforward. Thanks to the fact that sockets are also files, these same functions can often be used on sockets. For example, take a look at this following PHP script, `socket1.php`:

```
<?php
$connection = fsockopen ("www.linuxformat.co.uk", 80);
if ($connection) {
    fwrite($connection, "GET / HTTP/1.1\r\nHOST:
www.linuxformat.co.uk\r\n\r\n");
    while (!feof($connection)) {
        echo fread($connection,256);
    }
    fclose ($connection);
} else {
    print "Unable to connect!\n";
}
?>
```

That script opens a connection to **linuxformat.co.uk** on port 80, sends a HTTP request, then receives the web page back. Note

that **fsockopen()** is used rather than **fopen()**, because **fsockopen()** allows you to specify a port as the second parameter. **Fopen()**, on the other hand, opens a port based upon its best guess of what service you want. The advantage to this is that you can just **fopen()** **http://www.linuxformat.co.uk**, and PHP will spot that it's a HTTP connection and so open up port 80. The disadvantage is that you lose that extra control – **fsockopen()** lets you open up port 40391, port 23934, or port 39211. That is, you can connect to any port you want, which makes it a great deal more useful.

As **fopen()** automatically detects the type of connection you want, it also handles the protocol associated with it – the **fwrite()** line that sends a HTTP request is not needed when using **fopen()**, because **fopen()** sends it for you if you use port 80. Again, using **fsockopen()** requires more code, but you get much more flexibility as you can send whatever request you want.

The return value of the **fsockopen()** call is checked, and it will return false if the connection failed or a resource if everything went smoothly. The return value, on success, is on the face of it exactly the same as the return value sent back by a successful call to **fopen()**, and it can be used in the same way. Behind the scenes, of course, things are a little more complicated because the connection exists over the network, but this is all entirely transparent to us as PHP programmers.

To request a page from the server, we need to send a HTTP request. This is done using a HTTP GET request, as shown in the **fwrite()** line. Note that there are line breaks in there, so the string inside **fwrite()** is in fact two lines being sent. The HTTP standard is quite easy to read, but, like most protocols, is very strict. The first line requests `/`, the root file on the server. This is followed by **HTTP/1.1** to denote the kind of request being sent, then a carriage return `\r` and new line `\n`. The second line in the string specifies the hostname we're trying to access – this is crucial in **HTTP/1.1** because of the need to support name-based virtual hosting. The last part of a valid HTTP request is two sets of carriage return/new lines.

Once **fwrite()** is called, we should be able to start reading back the content, which is where the **feof()** loop comes in. The **feof()** function returns true if the file pointer passed to it still has more to be read – in this situation it will return true if there is more to be read from the web server. So, our loop says “while there is more to be read, read in 256 bytes” and print it out.

Finally, the connection is **fclose()**, which is standard clean-up procedure. HTTP junkies out will know that given the above code the server will actually automatically close the connection, but there's no harm making sure!

Go ahead and try the script yourself – you should see a printout of the source code from the LXF website like the example overpage. However, now that we've got full control over the HTTP connection, we can alter our request to specify that we only want to receive the headers of the page using a **HEAD** request rather than a GET request. Save this next script as `socket2.php` and try it out:

```
<?php
$connection = fsockopen ("www.linuxformat.co.uk", 80);
if ($connection) {
    fwrite($connection, "HEAD / HTTP/1.1\r\nHOST:
www.linuxformat.co.uk\r\n\r\n");
    while (!feof($connection)) {
        echo fread($connection,256);
    }
    fclose ($connection);
} else {
```

```
print "Unable to connect!\n";
}
?>
```

Note that the sole change is in the **fwrite()** line – GET has become **HEAD**, which instructs the server to only send the HTTP header data and not the full page. If you're checking to see whether a particular page exists or not, this would actually be the best place to start – simply check the string for a 404 error, and you're set.

A PHP server

While its clear that **fsockopen()** allows more flexibility than just **fopen()**, it is still quite simplistic when compared to what *can* be done with sockets. We've just covered the basic use of sockets in PHP, and the next step is to move onto the advanced, more flexible sockets that allow you to really create cool stuff with PHP. If you've not fully understood the code before, I strongly advise against continuing without having a read through of previous articles in this series, because things get a lot harder from this point onwards.

The five key functions to work with are **socket_create_listen()**, **socket_accept()**, **socket_write()**, **socket_read()**, and **socket_close()**. Combined, these sockets do much the same as the basic functions, except they also automatically handle *blocking* – pausing execution of the script until something important has happened. The advantage to this is that we can launch a PHP script and wait for things to happen, as opposed to making things happen. Sound like anything you recognise? That's right – a server!

Our new functions collectively handle everything we need to make a server in PHP. Precisely *what* we'll serve is beside the point, at least right now. The new functions create and instruct a socket to listen on a given port, receive a client connecting to the socket, write text out to the client, read text sent by the client, and close the socket. To get started with the advanced socket usage, you need to call **socket_create_listen()**, which takes a port number to listen on as its only parameter. This function creates a socket, binds it to the port specified in the parameter, and returns as a resource the socket it created or false if it failed. The socket resource returned from **socket_create_listen()** is used by other socket functions, so you'll almost certainly want to stash it away for the time being.

Moving on, once the socket is open, our server needs to accept a connection. This is done using the **socket_accept()** function, which takes the return value of **socket_create_listen()** as its only parameter, and returns a client connection – someone who connected to our port number, presumably looking for our server. **socket_accept()** operates by examining the queue of people waiting to be served, and taking the first client from there – if there are no clients waiting to be served, **socket_accept()** will wait until a client does become available, at which point it will return that. This action of waiting, usually called 'blocking', is crucial to our script running as a server – we're waiting to be asked for data, as opposed to actively requesting data.

Once we have a connection, we can crack on and communicate with it. **socket_write()** takes two parameters, which in order are the client to write to and the value you want to write – this data is then sent to the client in the order provided, as you'd expect. Its partner, **socket_read()**, also takes two parameters, which are the connection to read from, and the number of bytes to read. By using **socket_write()** and **socket_read()** together, you can fully interact with clients connecting to your socket.



« General to specific

Enough blathering, it's time to see the code – after all, you're not reading this article because of my hilarious jokes, right?

Servers do one thing: serve. What they serve is down to us, so we're going to pick a suitably easy thing to serve for our first server – to coin a word, we're going to make an 'uppercaseriser'!

```
<?php
// this socket number is random
$socket = @socket_create_listen("55555");
if (!$socket) {
    print "Failed to create socket!\n";
    exit;
}

while (true) {
    $client = socket_accept($socket);
    $welcome = "\nWelcome to the Magnificent Uppercasing
Engine.\nType '!exit' to close this connection, or type '!die' to
halt the server.\n";
    socket_write($client, $welcome);
    while (true) {
        $input = trim(socket_read($client, 256));
        if ($input == '!exit') {
            break;
        }
        if ($input == '!die') {
            socket_close($client);
            break 2;
        }
        $output = strtoupper($input) . "\n";
        socket_write($client, $output);
        echo $input\n";
    }
    socket_close($client);
}
socket_close($socket);
?>
```

Remember what I said earlier: using sockets in this manner is generally best done using the CLI SAPI, because it has no script timeout set. Run the script in the CLI SAPI, in a new terminal, fire up telnet and connect to port 55555 on localhost, like this:

```
telnet localhost 55555
```

That should launch your telnet program, which is useful for

forming simple connections to servers. All being well, you should receive the welcome message from the PHP server we just started it – try it out with a few sentences, then type !die to terminate the connection.

At this point you should have a basic server working that automatically uppercases text that you send to it. Yes, it's simple, but hopefully you should be confident using sockets properly now, which means we can move on.

Making a web server

Perhaps the simplest, useful server we can create using PHP is a web server, which, at its core, is basically an application that listens for requests on a given port (usually 80), and serves up the files that are asked for. Most web servers do a lot more than that, with common tasks being receiving POST data, handling file uploads, permissions, filters, and lots more. However, as creating such a script would require thousands of lines of code, we're just going to go for the lowest common denominator: serving files.

As we saw earlier, when looking at `fsocketopen()`, an average HTTP request looks like this:

```
GET /somefile.php HTTP/1.1
```

```
HOST: www.linuxformat.co.uk
```

Now, don't get me wrong on this point: adhering to standards is a very important thing, and I would never ordinarily recommend ignoring or contravening standards in programming scripts.

However, writing a server that properly parses and checks the entire HTTP request is out of the remit of this tutorial – you'll have to do that yourself. To serve up files, which is our goal, we just need to know which file was requested, and that's all in the second line. Therefore, to grab which file we need to send, we just check for the text after GET and before HTTP/1.1, and bin the rest. The simplest way to do this is to explode the request string as an array, and read in element #1 (arrays are zero-based, remember).

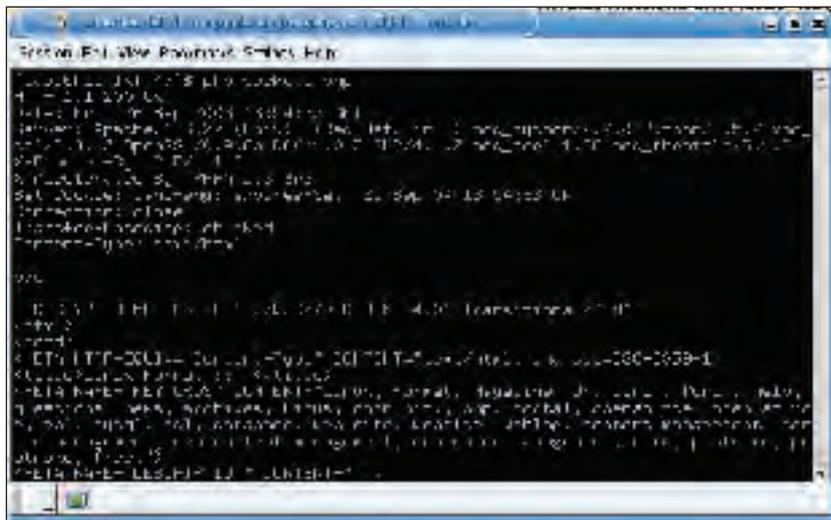
Each HTTP response needs to take a very specific format because HTTP splits its response into "header" and "content". If you don't take care to adhere to the standard when responding to a request, it's quite likely that clients will turn their nose up at what you send them. A properly formed request should contain several lines of header information followed by a couple of blank lines, then the actual HTML to be displayed. The header content should include the name of the file being served, its size, its type, the name of the server, and also the HTTP response type. The content part is the file itself, and is separated from the header by two sets of carriage return/line feeds.

HTTP response types are one line of text that contains the HTTP version being used (usually 1.0 or 1.1), followed by the HTTP status code of the response being sent and a small amount of text explaining the status. The content type is particularly important because it tells browsers what kind of file is being sent, upon which browsers usually base their decision on how to handle the file. Content types are defined as part of a standard known as Multipurpose Internet Mail Extensions (MIME) to define the exact type of data, for example the MIME type for HTML is 'text/html', and for GIF images it's 'image/gif'.

Now you've got a working understanding of how HTTP works, it's time to get onto the code. Note that I'm opening the socket on port 8000, as port 80 is often closed to non-root users.

```
<?php
$httpsock = @socket_create_listen("8000");
if (!$httpsock) {
    print "Socket creation failed!\n";
}
```

Julian, the LXF Art Editor, has banned us from showing screenshots of the LXF website – we wonder if he'll notice this roundabout pic...



```

exit;
}
while (1) {
    $client = socket_accept($httpsock);
    $input = trim(socket_read($client, 4096));
    $input = explode(" ", $input);
    $input = $input[1];
    $fileinfo = pathinfo($input);
    switch ($fileinfo['extension']) {
        default:
            $mime = "text/html";
    }
    if ($input == "/" ) {
        $input = "/index.html";
    }
    $input = ".$input";
    if (file_exists($input) && is_readable($input)) {
        echo "Serving $input\n";
        $contents = file_get_contents($input);
        $output = "HTTP/1.0 200 OK\r\nServer:
APatchyServer\r\nConnection: close\r\nContent-Type:
$mime\r\n\r\n$contents";
    } else {
        $contents = "The file you requested doesn't exist. Sorry!";
        $output = "HTTP/1.0 404 OBJECT NOT FOUND\r\nServer:
BabyHTTP\r\nConnection: close\r\nContent-Type:
text/html\r\n\r\n$contents";
    }
    socket_write($client, $output);
    socket_close($client);
}
socket_close($httpsock);
?>

```

Save that script as `webserver.php`, then start it up – you should be able to get to it through your web browser by visiting **http://localhost:8000**. You'll need to put a little content in there to make it work at its best – create a simple HTML file to get started, but note that anything other than HTML won't work currently.

The functions that will need explaining are `socket_read()`, `pathinfo()`, `file_get_contents()`. `Socket_read()` is basically the reverse of `socket_write()`, and reads data sent by the client into the return value. `Pathinfo()` reads the name of a file, and breaks it up into its component parts – filename, directory, and extension. The return value of `pathinfo()` is an associative array of these values, so we use `$fileinfo['extension']`. Finally, `file_get_contents()` opens the named file for reading, returns the complete contents of that file, then closes it again.

Advanced HTTP

There's a lot more you can do to our **BabyHTTP** to make it a little more powerful. For example, note that there's a switch statement to handle various MIME types, and only HTML is supported right now – if you add an entry in there for PNG pictures (MIME type 'image/png'), you could modify your sample HTML page and have it embed pictures in there. Sound tricky? We're not accepting any sort of keep-alive information with our server, which means we want each request as a whole new connection.

In this situation, a web browser will connect once, download the HTML, close the connection, read the HTML looking for subsequent files that are required (such as images, CSS files, Flash movies, etc), then make a separate connection for each one. To

MAKE YOUR MARK

Brainstorms 'R' Us

We're always looking out for ideas for new *Linux Format* PHP tutorials, and where better to look than to you, our readers? If, while reading past issues of *LXF's* PHP tutorials, you've thought "I wish they'd covered XYZ in more depth...", or "I really want to know how to use...", then now's the time to get your voice heard! Send an email to paul.hudson@futurenet.co.uk with your ideas – all the good suggestions will be covered in future issues. If you're short of ideas, you're certainly welcome to write in or post on the forums at www.linuxformat.co.uk with comments – we're passionate about improving the overall quality of our tutorials!

handle this we just need a new MIME type for each file we want to send because, essentially, the actual contents of the files being sent is irrelevant because `file_get_contents()` is binary-safe.

Once you've added support for a few picture types and perhaps also CSS files, try creating a new HTML file that includes images – it should all work fine. In fact, once your server accept requests for CSS files, pictures, and Flash movies, you should be able to serve up quite complicated pages – not bad for such a simple server!

Although adding support for a variety of files is likely to add the most immediate benefits, there are two other possibilities you might want to look into that will improve the server whilst also helping you learn a lot more about PHP. As you will have noticed if you tried creating a complicated HTML page, our server only handles one connection at a time, even if there are more waiting. This is because our server exists in only one execution thread – PHP executes the code linearly, one connection at a time.

While this is perfectly fine for fairly simple servers, it would be better if the server created a new process to handle each request as soon as it came in, and this is accomplished in the PHP world by using the `pcntl_fork()` function – something a little too complex to cover in this article, but definitely something you should read up on if you want to improve your server.

The second potential improvement is having the ability to serve up compressed content, which is a little easier to do. The key is to look for the **accept encoding** line in the HTTP request coming in from your client, as it will tell you what kind of compression the client can accept. *Konqueror*, for example, reports back **x-gzip**, **x-deflate**, **gzip**, **deflate**, **identity**, of which we're interested in **gzip**, as it shows it can receive gzipped content – give it a try!

Conclusion

Sockets are a fun and interesting part of PHP, more so because they are one of the least-explored parts. As we've seen, it's actually quite easy to set up a simple server to handle our own tasks, although perhaps uppercasing user input is a little too easy – other possibilities include an encryption server, or a database querying server, etc. We've also looked at how it's actually quite easy to move from a simple server to a basic web server – granted our server isn't a patch (sorry!) on *Apache*, but, once you program in the ability to handle other media types, it can handle fairly complex stuff.

If you're thinking that sockets are only useful with the CLI SAPI, you'd be quite wrong – using sockets you can connect to all sorts of services on the Internet, such as the O&A currency system (www.oanda.com), or Unix time servers. There is quite literally a massive range of options available, of which HTTP is just one – experiment! 

NEXT MONTH

While there's nothing more we love to do than write PHP code, sometimes it's easier to re-use the hard work of everyone else. As such, we'll be looking over the **PEAR** code repository, and examining what it can do for us...



EVERYONE SHOULD HAVE A LINUX SERVER

SERVER SCHOOL: Build a basic firewall

PART 3 A firewall is an essential requirement in these days of worms and other Internet nasties. **Jono Bacon** introduces the core technologies and concepts behind them.

Recently, the Internet has become a pretty dangerous place. The advent of worms and viruses targeting our computers is forcing us to make additional steps in keeping our systems safe. Although every technology can be fundamentally comprised if a cracker is determined enough and has a high enough level of technical proficiency, there is one basic piece of protection that we can use to protect our machines and networks: firewalls.

For those who have never used a firewall before, a firewall is a hardware and software protection device that is put between your computer and the Internet. It is best to imagine a firewall as a literal concrete wall between yourself and the Internet. This concrete wall will block anything that wants to talk to your computers, apart from specific types of connection. These allowed connections conceptually speak through little holes in the firewall so that the bulk of the wall is still defending the machine.

Setting up the hardware

Although you can create a firewall in software on your normal computer, this method has one distinct disadvantage – if an intruder manages to bypass your firewall, they will have access to your normal computer – with all our sensitive information on there, this is obviously something that users everywhere will want

to avoid. A better solution is to use a dedicated computer as a firewall. This method has the opposite benefit that if the intruder gains access, they can only access that machine which will have nothing interesting or useful on it. If the intruder then wants to access your main machine on the network, they will need to bypass another level of security. Two levels of security and isolation of the firewall are generally considered good things, as long as the security does not unreasonably infringe on your normal usage of your computers.

Using a dedicated machine may sound like an expensive proposition, but it really doesn't have to be. The processing requirements for a firewall computer are very low, and you can reasonably run a firewall on a very low spec 486/Pentium machine with 64MB RAM. Disk space is also a low requirement, and the only thing that takes up space is system logs of firewall access and denied requests. As such, if you go to a computer fair, you can often pick up a pre-owned machine such as this for very little. The only additional requirements that you will need is dependent on your type of modem. If you are using an Ethernet cable modem, you will need two network cards (one for the cable modem and one for the main network), if you are using a USB/Serial modem, you will need the Serial/USB port and an Ethernet card to connect to the network.

Introducing iptables

The software side of setting up a firewall is handled by *iptables* which is part of the Netfilter project. *iptables* came about as a replacement to *ipchains*; a software facility to construct chains of rules that determine how traffic behaves on a network. Although *ipchains* worked well, a replacement was proposed and *iptables* was born, with an *ipchains* compatibility mode included.

iptables is a packet-filtering technology. A packet is the smallest form of network traffic, and is similar in concept to an envelope. On each packet there is a destination address, a source address, the contents of the packet (the data) and some additional information. When you use *iptables*, we will check these additional bits of information and filter out those packets that we do not want to accept in our network. Not only is *iptables* useful at checking packets to see if we want to accept them, but it can be used to send packets to different parts of the computer or network dependent on what they are for. As an example, if you have a web server that you would like to make available publicly, this server would possibly sit inside your network. This means that when someone connects to your network, instead of requesting a web connection on your firewall machine, you want another machine to handle the connection. This requirement is called port forwarding, and *iptables* can help with examples such as this.

Installing iptables

iptables is available right within the kernel. This can be good or bad for you, dependent on how comfortable you are compiling a kernel. If you have never compiled a kernel before, don't worry, as there is plenty of information on the Internet to help you do it. To give you a head start, here is the series of command that you enter to compile your kernel. We assume that you have downloaded the latest kernel and put it in `/usr/src`.

```
gunzip 2.x.x-source.tar.gz
```

or

```
bunzip2 -d 2.x.x-source.bz2
```

Then

```
tar xvf 2.x.x-source.tar
```

make config, **make menuconfig** or **make xconfig** (Select the features you would like in the kernel)

```
make dep
```

```
make bzImage
```

```
make modules
```

```
make modules_install
```

Copy `/usr/src/linux/arch/i386/boot/bzImage` to `/boot`

Edit `lilo.conf` and then run `/sbin/lilo`

If this is new to you, it is recommended that you consult the Kernel HOWTO that is available at www.tldp.org/HOWTO/Kernel-HOWTO/; it will explain each of these steps in detail.

When you are configuring your kernel (with the **make config**, **make menuconfig** or **make xconfig** step), you will need to ensure that you include the right features to use *iptables*. Much of what is required is highly dependent on what you will be using *iptables* for, but the following options should get a basic system up and running:

NETWORKING OPTIONS

- Packet Socket
- Network packet filtering
- Network packet filtering debugging
- Socket Filtering

IP NETFILTER

- Connection Tracking
- *iptables* Support
- Limit match support
- MAC address match support
- Packet type match support
- Netfilter MARK match support
- TOS match support

There are of course many other *iptables* options available, and these options may be required for more exotic uses of *iptables*.

Using iptables

iptables makes use of a number of 'rules'. These rules specify conditions that can be applied to the traffic that is coming into your computer and deal with it accordingly. These rules are specified in either the command line or as part of a more expansive shell script. Writing the rules is relatively simple and merely involves knowing what traffic you are managing with and how you want to deal with it. Although a rule is typically a single line in length, these rules can be spread over multiple lines by appending the rules together.

Before we write our first rule, it is important to remember that you will need to be able to test your rules sufficiently to ensure that they work. As an example, if you are wanting to block all mail related traffic accessing your machine, you will need to ensure that you can send some mail traffic to test the rule. This is an area in which writing rules can be a little tricky, particularly if you are dealing with traffic for services that you do not have access to. Luckily though, most intended uses for firewall rules are for blocking everything apart from specific services, so you do not necessarily need to have these services available.

To get us started, I will work on a simple rule that not only demonstrates how to write an *iptables* rule, but can be used on virtually any Linux machine. This rule will block traffic from the localhost loopback IP address (127.0.0.1) that refers to the current machine that Linux is running on. You will need to have support for the loopback address compiled into your kernel, but most Linux distributions and kernels have this support as standard. To test, you can use the *ping* tool to send packets to the following address:

```
ping 127.0.0.1
```

You should receive something such as this as a response:

```
64 bytes from 127.0.0.1: icmp_seq=0 ttl=64 time=0.2 ms
```

```
64 bytes from 127.0.0.1: icmp_seq=1 ttl=64 time=0.1 ms
```

```
64 bytes from 127.0.0.1: icmp_seq=2 ttl=64 time=0.1 ms
```

When you press Ctrl-C to stop the ping, you will then see:

```
3 packets transmitted, 3 packets received, 0% packet loss
```

This *ping* tool sends a packet to 127.0.0.1 and reports back how long it took to get there. We can also see there is a 0% packet loss which means that all packets got to 127.0.0.1 successfully. We can use *iptables* to block these packets if required, and we know some important information already. First, we know that the address to block is 127.0.0.1; second, we know we want to drop this traffic; and thirdly, we are using *ping* – which means that the type of traffic is ICMP (this may not be obvious, but most services have their packet type available in the documentation). Using this information we can construct the following rule:

```
iptables -A INPUT -s 127.0.0.1 -p icmp -j DROP
```

You will need to type this rule in as root, and if successful, you



◀◀ will be returned silently to the prompt. The rule states that using **iptables**, we will append a new rule to the system **-A** using the INPUT chain referring to packets from 127.0.0.1 (**-s 127.0.0.1**) using the ICMP protocol (**-p ICMP**) should be dropped (**-j DROP**). You can see that this command reads from left to right in a relatively simple manner, but can be confusing at first due to the number of different options. When you now run the ping command you will get 100% packet loss due to the rule dropping the packets.

To bring back the ping-ability (yes, I know it's not a word, but it sounds cool) of your system, you can delete the rule with:

```
iptables -D INPUT 1
```

This command refers to the first rule on the system. You can get a listing of all of these rules by running:

```
iptables --list
```

Exploring iptables further

With our basic knowledge of *iptables*, we will now take a look at *iptables* further and some of the additional facilities to filter your traffic. As with any technology, these are the basics, and for more detailed information, you should check the man page for *iptables* and other documentation on the Internet.

The first objective we may have for our firewall is to block certain types of request coming into our machines. A typical example of this is blocking the insecure Telnet service. Telnet can be used for remote access, but is unencrypted and *ssh* offers a far better solution. *iptables* can block telnet traffic by specifying it as a port service in a rule. As well as blocking this access, we also need to specify where the access is originating from, which is usually your modem or network card. This hardware is referred to as an Interface in Linux, and you can see the currently working interfaces by typing:

```
ifconfig
```

We will assume that you need to block all traffic coming in from your Ethernet card (**eth0** is typically the name of the interface). This could be blocked with:

```
iptables -A INPUT -p tcp --destination-port telnet -i eth0 -j DROP
```

This command uses the **--destination-port** switch to specify telnet as a port service on the **eth0** interface (with **-i** switch). Although using the **eth0** interface is useful to handle all traffic, you may want to restrict traffic to certain hosts. These hosts can be specified by either using IP addresses or hostnames and the **-s** (source) and **-d** (destination) switches. As an example, let us

assume that you want to restrict traffic from a single host, you could use the command:

```
iptables -A INPUT -s 192.168.0.100 -j DROP
```

A useful addition to this use of IP address and hostnames is the **!** symbol. This means 'not' in the world of *iptables*, and you can use it to specify all traffic other than the host you specify. As an example, you could block all traffic other than that from 192.168.0.5:

```
iptables -A INPUT -s ! 192.168.0.5 -j DROP
```

Getting all GUI about firewalls

At this point in our exploration of firewalling, we have taken a look at *iptables*; the technology that drives a firewall. The reason we have taken a direct look at this technology is so you have a basic foundation knowledge of what is going on behind the scenes. For most practical firewalling, you may want to make use of a graphical interface to *iptables* that can save a lot of time and error when building your custom firewall. There are many graphical firewall utilities available, and each of them applies its own method of building your firewall rules. Of all of the tools I have tested over the last few years, one of the most impressive that I have used is *Guarddog*.

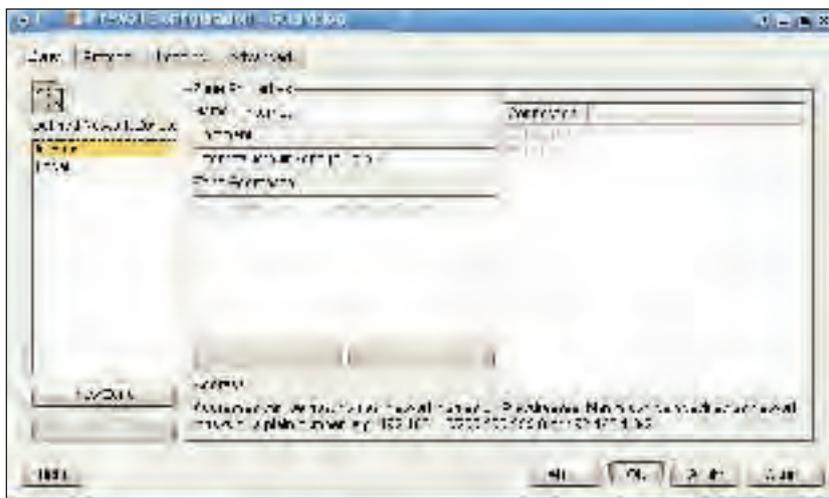
Guarddog is a KDE-based firewall interface that is not only expansive in what it can actually do, but is impressively simple in its use. The application is available for many distributions and the homepage is at www.simonzone.com/software/guarddog/. The only requirements for *Guarddog* are the KDE libraries and a 2.4 kernel (which you will need to use *iptables* anyway!)

Guarddog has a simple interface that is spread over a number of tabs; by going through each tab the firewall can be configured. Fig1 shows the main window of *Guarddog*, and on the left hand side is a number of different zones. Each zone can be created to refer to a general network that you will access. This already includes Internet and Local, but you may want to add additional networks with the New Zone button. The middle part of the interface refers to any IP addresses that may be specified. For the purposes of this article we will focus on adding a firewall to protect us from the general Internet. The built-in Internet zone has no specific IP addresses specified so *Guarddog* will apply the firewall to all hosts.

If you click on the second tab; Protocol, you will be able to select which services are allowed and blocked. The main Zone Properties area of this tab displays a number of categories of network services. When you expand this tree of categories, you will see a list of services and a checkbox by each one. By default all of the services will not be checked and this means all of them are blocked. You will now need to select those services that you would like to allow through the firewall by ticking the box. If you click one of the checkboxes twice you will be able to place a cross in the checkbox that will outright reject those requests as opposed to simply blocking them. Each service that can be configured has a short description of the service and its potential risk to the left of the interface. You can see some of these services configured in Fig2 on the opposite page.

The third tab in *Guarddog* is the Logging page. This section allows you to configure not only what is logged, but how often the logs are updated. It is highly recommended that you log as much information as possible, and if you are unsure log everything. The main areas to log are blocked and rejected packets. If you log these attempts to access your machine, you will be able to see how well your firewall is preventing access from unwanted and possibly malicious traffic.

Fig1 The opening tab in the *Guarddog* interface.



DEDICATED FIREWALLING WITH SMOOTHWALL

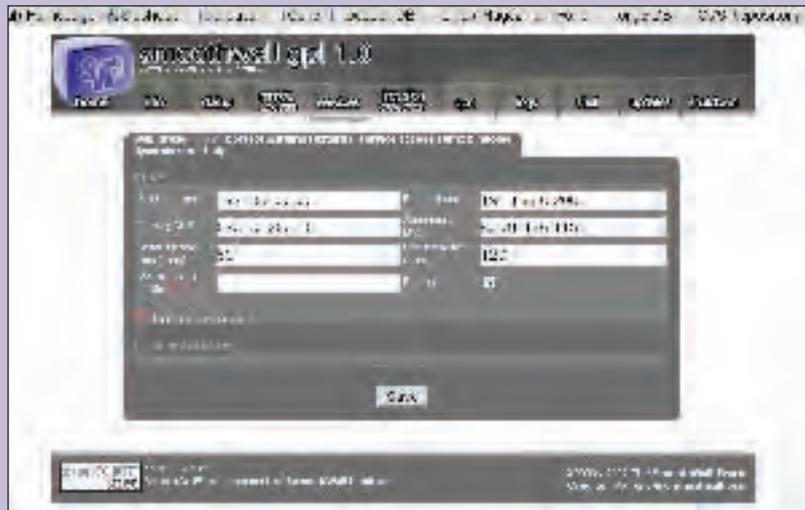
Simple to set up and install

Although *iptables* is a powerful and flexible part of the Linux kernel, it can be a little daunting setting all of the software up properly and ensuring that everything is up to date. To fulfill this need, organisations and companies have formed to develop special Linux distributions dedicated for firewalling. These distributions are not intended to be used as a desktop distribution, but are usually used on old low-spec PCs that are dedicated firewalls.

Of these distributions, one of the most popular is *Smoothwall*. This product supplies a complete web based firewall with DHCP server, intrusion detection system, port forwarding,

external services, dynamic IP support, VPN and many other features. Not only does *Smoothwall* provide most features you would need for a firewall, it is simple to set up and install. You can download *Smoothwall* from the website at www.smoothwall.org/ as an ISO image that can be burned to a CD. When the system is installed you can keep it up to date with regular software patches that are released to fix published security flaws and problems.

This always means that you can be up-to-date and secure without having to source the fixes yourself; they will be listed on the Updates page in the *Smoothwall* web interface.



The *Smoothwall* web interface is both attractive and very straightforward to use.

The remaining area in *Guarddog* that we have not mentioned yet is the Advanced tab. In this page there are variety of options for configuring *Guarddog* and how it works. The first box to note is the Disable Firewall checkbox. You can use this box to temporarily disable the firewall without having to re-implement it every time. The 'Show advanced protocol help' box will show additional details in the Protocol tab page for each service. This can be useful for more technical users. Another function on this page is the ability to add custom protocols that you have set up on your machines. As an example, if you have created a number of *Apache* virtual servers, you may have given them different port numbers such as 8001 and 8002. You can add these ports with the New Protocol button if they need to be accessed outside of the firewall.

Summary

In this article we have tried to prepare a good grounding in the issues faced when setting up a firewall. Not only have we discussed the hardware and software issues for a firewall, but we have also looked at the underlying *iptables* technology and a graphical interface to it. Combining knowledge of each of these areas will help you to build a more secure firewall.

There is no doubt that firewalling is a complex and sprawling subject, and there is much more to learn to secure your system successfully. It is recommended that you look into other security technologies such as intrusion detection systems, advanced logging systems and encryption to ensure that your network and computers are not comprised when connected to the general Internet and other computers. [LXF](#)

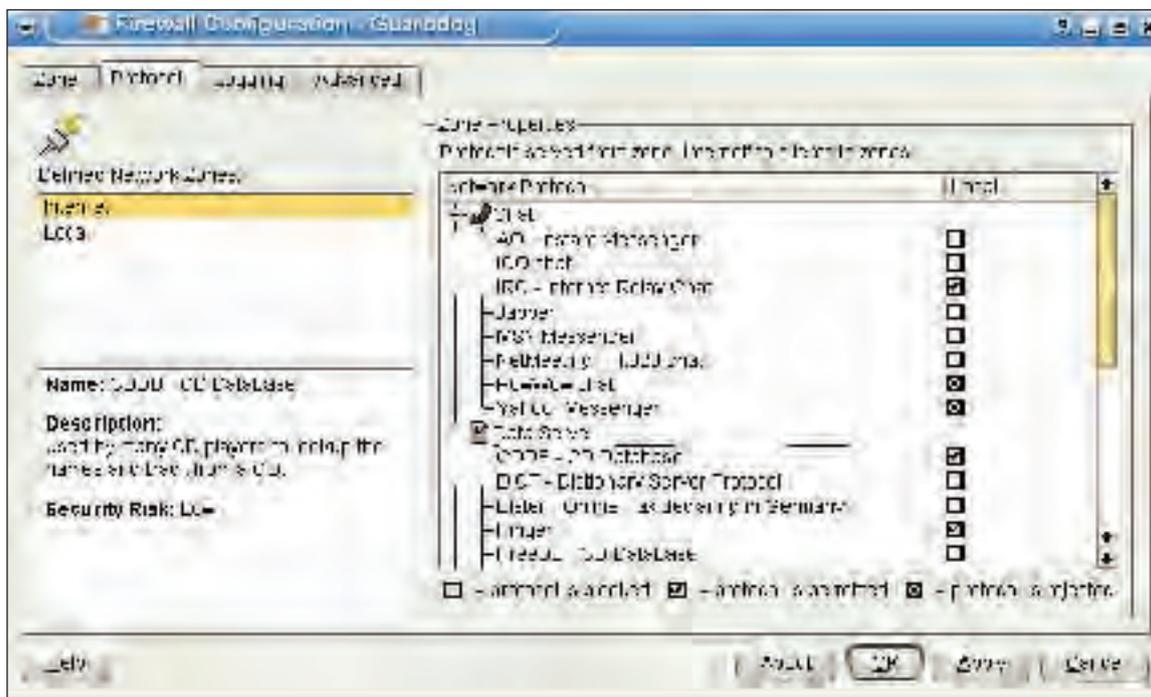


Fig2 Selecting services that can be allowed, blocked or rejected.

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



Mandrake 8 to 9

I have been suffering from a very irritating problem that was stopping me from using the latest Mandrake distribution.

I thought it was a problem with my hardware not being compatible with the new version (the Network card is a DLink DE530). The old version (8) works fine all Networking to our Broadband detects and works fist time, but 9.1 appears to fail.

Thanks to the letter *LXF43* page 14 regarding a similar type of problem with SUSE, I ran a quick test by disabling ACPI (Control of the first 6 PCI IRQs) and all was well with MDK 9.1 – the Network was working perfectly.

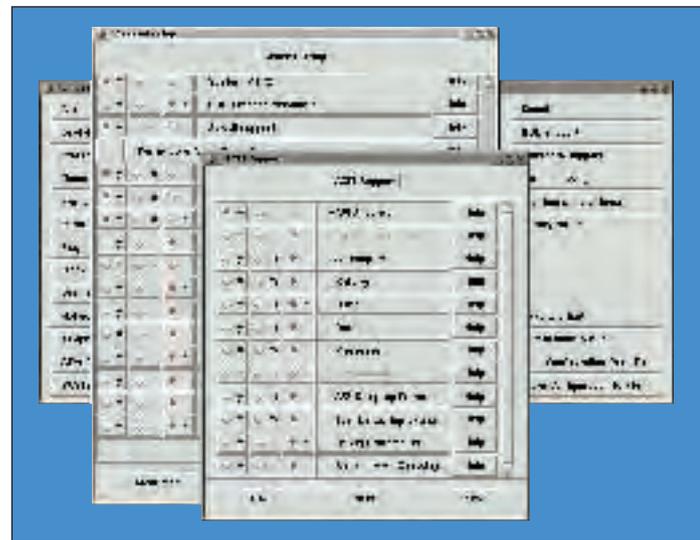
Now for the down side, as stated in the *ABIT KD7-E User Manual*; if you install Windows XP or 2000 with ACPI Enabled and then Disable it, Windows will no longer work until it is re-enabled. At least I now know how to work around this problem (disable ACPI and re-install Windows).

Whatever has caused this problem then it must be some change in the code between 8 and 9, does anyone know why? And is it likely to be fixed?

Trevor Williams, via email

As the Linux kernel changed between Mandrake 8 and 9.1, it's likely that some incompatibilities have been introduced into the distribution. Since disabling ACPI fixes the problem, it appears to be a known issue with Mandrake, although we're unsure exactly how the Mandrake developers are approaching it.

You may want to look at upgrading the kernel that is distributed with 9.1 to a more recent build, as it may solve the immediate issue. We supply up-to-date versions of the kernel on the *LXF* coverdiscs each month.



ACPI is an alternative to APM on many systems, and provides far more power management features.

APT

I am still very wet behind the ears when it comes to Linux, I would very much like not to be, so I am reading your mag with interest especially the APT. I followed your instructions and – hey presto – it all worked. So my question is: why are these sort of instructions not on the disk for the software that is on there? For example, I thought to have *Xine* on my PC would be nice so that I can watch *Lord Of The Rings: The Two Towers*, so I go to the relevant folder and I am beset with .tar files, SRPMs & RPMs.

I think I will use *apt-get* on the RPM files but every one has a 'cannot find files' error. Didn't even work with the SRPM files. I found your instructions on how to do .tar files and after doing three of them, the last one did not have any errors. I thought "Great, I can now watch a DVD." I went to the main menu on the desktop and looked for *Xine* but it was not to be found. I plodded through the folders not knowing where to go in the vain

hope of finding something that would say "double-click me", but no... Hold on I think, if I put a DVD in, then it should start up straight away, shouldn't it? I got a DVD down and with mounting excitement I placed it on the tray watched it go in and strained to hear the drive whirr away, but no. So crestfallen, I write to ask if it at all possible to give some sort of indication as to what file to use because I haven't got a clue. I have been using Windows for quite some time, but I would like to use Linux more – a clue as to the equivalent to the .exe file would be good. Similarly I couldn't get *Realsoft3D* to work either.

Stephen Jessop, via email

Taking your *apt-get* question first. Before you can start using it, you need to build a 'sources.list' file, based upon your distribution and specific version. This will allow *apt* to locate the correct files and download them for you. Once 'sources.list' is updated, one can do an **apt-get update** and download the updated package list for the

distribution. You should then be able to 'apt-get install' whatever you need. For RedHat users, a source of sources.list is at <http://apt-rpm.tuxfamily.org/>.

With *Xine*, if you install it from source, chances are that it's not going to install any menu items or desktop links for you to easily access it. The binary itself lives at `/usr/local/bin/xine`, so by running that from a terminal will allow you to access the DVD drive and play your movie DVDs.

Building from source code does not integrate the software with your system in the same way the RPMs do, as it is not distribution specific, so it's always a good idea to use packages, rather than attempting to build it from source unless you really need to. <http://rpmfind.net/> is a great resource for anyone looking for RPMs for software. Certainly 'apt-get' solves this issue, as it will automatically download RPMs and their dependencies.

MDK oops!

I am having a real nightmare regarding the HP CDwriter+ 9100c on My Dell Optiplex GX110 running Mandrake 9.0 . Every time I try to mount any CD, it just refuses to do that. The drive spins, the LED indicates the activity, IDE-controller LED is on. However, after half a minute ,the system replies it cannot mount the drive. I am using *ide-scsi* module so I can use *xcdroast* to burn CDs.

I had serious problems when I run 2.4.19 kernel – every time this situation occurred I got kernel oops. Then I tried it with vanilla 2.4.21 first, then with 2.4.21-ck1 (Con Kolivas' production patches) – same oops (bug in *ide-scsi.c* line 512, unfortunately I haven't got the oops printouts).

Now I am running 2.4.22-ck1 and fortunately I am not getting oopses, but still cannot mount any



The *ide-scsi* system is required to use a CD writer with Linux, but if the underlying IDE devices have issues, then it can cause confusion. >>

A QUICK REFERENCE TO: RPM

RPM is the package manager used by many distributions, including Red Hat, Mandrake and SUSE. Rather than having to compile everything from the source code, RPMs give the opportunity to install pre-compiled binaries of packages you need, ensuring that any dependencies, for libraries or third-party executables, are met before installation.

Installing a RPM downloaded from the Internet is done with the **-U** switch;

```
# rpm -U rpm-python-4.0.2-6x.i386.rpm
```

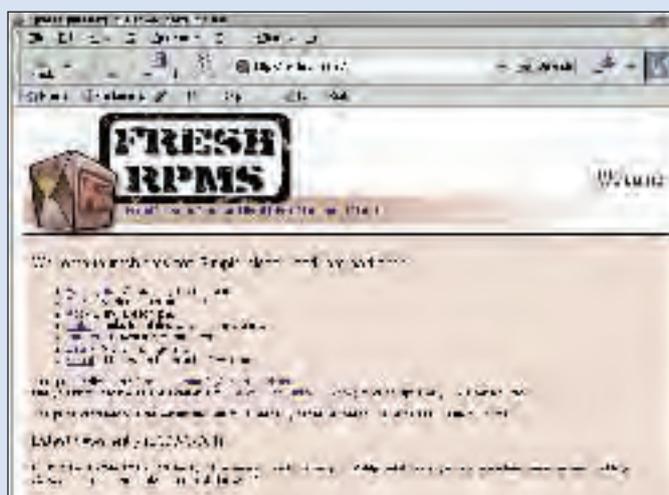
Before RPM will attempt to install the package, it will check existing packages which we have installed, and check that any required dependencies have been met. In this case, it complains;

error: failed dependencies:

libbz2.so.1 is needed by rpm-python-4.0.2-6x

libdb-3.1.so is needed by rpm-python-4.0.2-6x

In order to have this RPM install correctly, we need to install the package which provides *libbz2.so.0*, which is *bzip2*, and one which provides *libdb-3.1.so*.



Frustrated with dependencies on Red Hat systems? *apt-rpm* allows you to enjoy the ease of packet installation usually only found on Debian.

Once these two dependencies have been met, the original RPM will install correctly.

Of course, before we install an RPM, it's useful to know what it is. Performing queries on RPMs is done using the **-q** switch, and it's useful to know what dependencies are provided for by installing a particular RPM, which is done with **--provides**:

```
# rpm -qp --provides bzip2-1.0.1-3.i386.rpm
```

libbz2.so.1

bzip2 = 1.0.1-3

However, there are times when we know that dependencies have been met, either by a broken RPM, or by installing something from a source distribution, and we can have rpm ignore the dependencies of an RPM, and install it anyway:

```
# rpm -U --nodeps rpm-python-4.0.2-6x.i386.rpm
```

Naturally, if you really have not satisfied the dependencies for the

package, then it's not going to work at all, so having RPM ignore what it requires usually causes more problems than it cures.

While we can build from source, and install it without telling RPM about it, it is possible to build from the source code, create a RPM, then install it so everything behaves itself. Generally with RPM distributions, there are also *src.rpm* files, which are RPMs containing the source code for the particular package. The quickest way to build an RPM from the *src.rpm* is with **--rebuild**:

```
# rpm --rebuild bzip2-1.0.1-3.src.rpm
```

If the *src.rpm* does not exist, there is usually a *package.spec* file within the standard source tarball, although even if one does not exist, it's not a difficult task to create one by copying it from another package and adjust it appropriately. Building an RPM from the spec file is very simple;

```
# cd /usr/src/package-1.0.1
```

```
# rpm -bb package.spec
```

This will eventually produce a binary RPM in your `/usr/src/RedHat` directory, which can be installed like any other RPM.

FREQUENTLY ASKED QUESTIONS SPAMASSASSIN

FAQ WHAT IS SPAMASSASSIN AND WHY SHOULD I BE USING IT?

SpamAssassin is a mail-filtering tool specifically designed to identify spam. With a huge set of filters to check messages against, *SpamAssassin* can be integrated into an existing mail system to minimise the amount of unwanted email seen on a network.

SpamAssassin is a great way to reduce the amount of junk in a mailbox, limit excessive bandwidth on a network by blocking mail at an upstream mail server, and for businesses, improve productivity by avoiding having employees wade through hundreds of emails each day.

FAQ WHERE CAN I GET SPAMASSASSIN?

Most distributions have packages of *SpamAssassin*, allowing for a very quick and simple installation. Debian users can simply ‘

```
apt-get install SpamAssassin
```

and have it install all by itself.

As *SpamAssassin* is written in Perl, it is also available via CPAN. We can run the CPAN client with:



If you're seeing too much spam in your mailbox, *SpamAssassin* can be integrated with nearly all mail servers

```
# perl -MCPAN -e shell
```

Then, to install *SpamAssassin* you simply do:

```
install Mail::Spamassassin
```

LXF43 featured *SpamAssassin* on its coverdiscs – order your backissues from page 99 if you missed it.

FAQ HOW DO I ACTUALLY RUN SPAMASSASSIN?

SpamAssassin can be run either as a standalone process, or in a client/server configuration. For larger systems, running a server process makes more sense, as one does not have startup all of the Perl modules required for processing the messages for each delivery. The server also allows a single box to be dedicated to process messages to filter out spam,

then return the result back to a mail server, making it particular effective in a cluster configuration.

Most packages of *SpamAssassin* will provide an *init.d* script to run *spamd* at boot time, although some configuration may be required, as at least with Debian, *spamd* does not run as default at boot time.

FAQ HOW DO I FILTER MY MAIL WITH SPAMASSASSIN

Filter depends upon mail server specifics, although generally one can simply use *procmail* on a per-user basis to scan and filter mail. Minor reconfiguration may be required to have *procmail* handle your mail delivery, so check the mail server settings first.

It is also possible to handle *procmail* via the ‘forward’ file:

```
"IFS=' '&& exec /usr/bin/procmail -f - || exit 75 #user"
```

A sample *.procmailrc* would look something similar to the following:

```
DEFAULT=$HOME/mbox
```

```
MAIL=$HOME/mail/
```

◀ CD (not mentioning burning). What is even more mind-boggling is that when I am not using *scsi-ide* module to drive the HP CDwriter, it works fine! However, I can't burn CDs in Linux, or can I?

I tried to search all possible corners of Internet to find the help but with no success. I am aware that it might be an issue with the combination of disks connected to *ide* controller or loose or broken cables. I tried all the combinations, I changed the cables a couple of times, no success.

Posting to the forum The LXF online community

Not only do our popular forums at www.linuxformat.co.uk have sections dedicated to your technical queries, hardware, programming languages and general help; but also there's always a lively discussion going on!

I haven't got these problems when running Windows 98SE on the same machine – which drives me mad, as I know that the Linux is well capable to do the same job even better.

There is another slight niggle – Windows 98 can shutdown and power off Dell GX110 without any problems. However, Linux can't. I tried both options in the APM settings in kernel – using the real mode to power off or not. Still, after receiving the ‘Power down’ message the machine remains up and running.

To provide you with the crucial information, I am including all the printouts I could get hold of.

By the way, thanks for the great mag, I've been reading it from the issue 1 and thanks to your expertise we are now running a small lab with five dual-boot machines (Win2K / Mandrake 9.0), *Samba* server as PDC (roaming profiles) and *Smoothwall* to protect us from

outside. And we are thoroughly enjoying it

Petr Voles, Bournemouth University
Something looks very unhappy when you access the CD writer:

```
scsi : aborting command due to
timeout : pid 19, scsi1, channel 0, id
0, lun 0 Read TOC 00 00 00 00 00
00 00 0c 00
hdc: irq timeout: status=0xd0 { Busy }
hdc: ATAPI reset complete
hdc: irq timeout: status=0x80 { Busy }
hdc: ATAPI reset complete
hdc: irq timeout: status=0x80 { Busy }
scsi1 channel 0 : resetting for
second half of retries.
SCSI bus is being reset for host 1
channel 0.
attempt to access beyond end of
device
0b:00: rw=0, want=34, limit=2
isofs read super: bread failed,
dev=0b:00, iso blknum=16,
block=16
hdc: drive cmd: status=0x51 {
DriveReady SeekComplete Error }
```

```
hdc: drive_cmd: error=0x04
```

We'd suggest that you try to mount the device using the standard *ide-cdrom* module, rather than *ide-scsi*, just to ensure that the device works correctly and the media is good. The fact that it's saying it's trying to access beyond the end of the device would imply that either it is reading the media incorrectly, the media is empty or damaged, or simply there is something very wrong with your kernel. If you can debug this without having the SCSI system in the mix, it may help.

On the power front, it may be that your system uses ACPI for power management, rather than APM, as is the case with most modern systems. Usually it's more of a magic trick to get Linux to power down the system with APM, unfortunately, although some have success with it.

F*@\$% firewall!

I have a problem with my firewall, that I cannot fix. I run a small

```
:Ofw: spamassassin.lock
| /usr/bin/spamassassin
```

```
:0:
* ^X-Spam-Status: Yes
$MAIL/spam
```

FAQ HOW CAN I FINE-TUNE SPAMASSASSIN TO IMPROVE SPAM DETECTION ON MY SYSTEM?

SpamAssassin supports a number of different options, such as the use of DNS Real-time Blackhole Lists, checking against *Razor*, which contains a massive collection of known spam messages. *Razor* requires a third-party tool from <http://razor.sourceforge.net/>, which communicates with the service over the Internet.

There is also *Distributed Checksum Clearinghouse* (www.rhyolite.com/anti-spam/dcc/) and *Pyzor* (<http://pyzor.sourceforge.net/>) which perform the same function as *Razor*, but are implemented on different servers. Of course, having more valuable spam to compare incoming messages to massively improves the likelihood of known spam email being caught.

network of 3 computers:

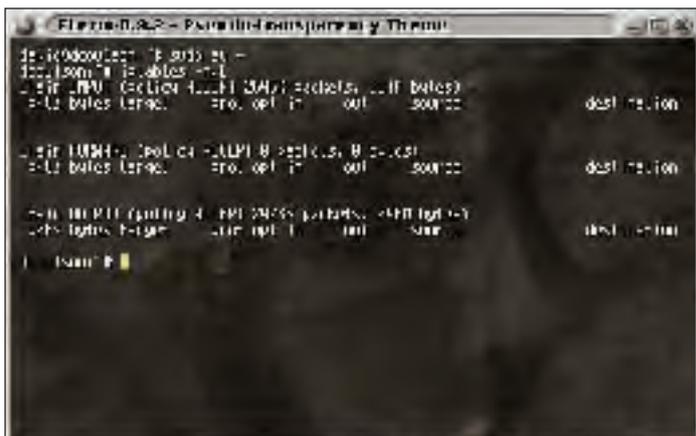
A is a MDK 9.0, the connection to internet

B is a Win98 (for my girlfriend)

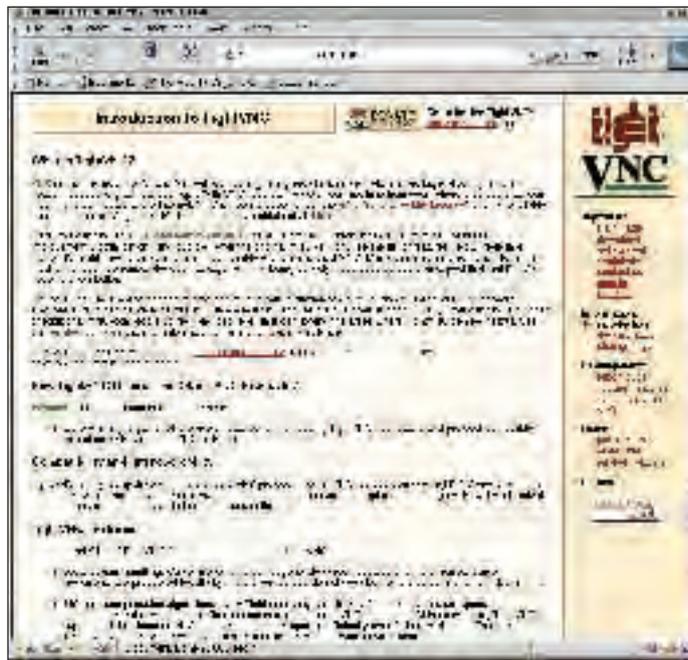
C is a MDK 9.1 (for me)

I have in `/etc/rc.local` a line for NAT:

```
iptables -t nat -A POSTROUTING -o eth1 -j MASQUERADE
```



Most distributions have a very basic firewall installed as standard, which can cause problems with manual configurations.



VNC allows Linux users to access Windows desktops, and vice versa, although the numbering of the displays can become a little confusing.

Every time I reboot **A** (The server, once every 3 months) I get:

```
ping www.yahoo.com (from A, B and C)
unknown host
```

so I go to **mcc** and in security, disable the firewall, run the NAT command `ping www.yahoo.com` (from **A**) get answers, from **B** and **C** – forbidden access. So I go back to **mcc** and restart the firewall, run my NAT command and then everything is working again. This is the firewall from **mcc**, and I have no idea what file is touched by it...

From the LXF forums

An 'unknown host' error comes from a failure to perform a DNS request, which should be reasonably easy to track down. Either your Internet

connection is down, your DNS server is down or unreachable, or `/etc/resolv.conf` is misconfigured.

We would disable the Mandrake firewall totally, to avoid any duplicate efforts with your `iptables` command. If you could provide us with the output from `iptables -nvl` and `iptables -t nat -nvl` both before it works, and while it is working, then it may prove useful in diagnosing the issue at hand. It may be that the firewall is blocking something, IP forwarding is not enabled, or services are not running correctly. Of course, if DNS does not work from the local host, then it's more than likely a firewall issue.

TightVNC

Having given up trying to get SSH working between Red Hat 8.0 and Windows XP (without using *Cygwin*) I downloaded the latest version of VNC for both Linux and XP but receive the error 'Connection Refused' or 'Failed to connect to server' whichever way I try to connect.

I have no firewall on the XP side and the Red Hat side is a pretty standard install. Any ideas as to why I get these messages? I used VNC with Red Hat 7.2 and Win2K a couple of years ago without problem, the VNC version is 3.3.7 and XP professional SP1. Any help would be appreciated!

Pete Kane, from the LXF forums

VNC requires a server to be running on the remote system before you can connect to it. With Windows, this just shares the existing desktop with any client that connects to the server. Clearly, if you can't connect to it, then it is not running. The server needs to be manually started from the 'Start' menu. Once the server is running, it will appear on the system tray over on the right side with a tiny VNC icon, so you know it's working. You should just be able to point the client at the IP address of the system and it will connect and ask for authentication information.

VNC on Unix is a little different, as it actually starts up a separate X session with its own display. It is run using the `vncserver` script, which will allocate an X display, generally `:1` for the first VNC session. The client needs to reference the specific X session number, as by default it will try `:0`, which in the case of a box running a standard X server, will not be capable of handling a VNC client. You will need to ensure that there is no `ipchains` or `iptables` firewall running on the Linux box, otherwise you may be blocking connections to the ports that VNC uses for client connectivity.

Submission advice

We are happy to answer all sorts of Linux related questions. If we don't know the answer, we'll find out for you! But in order to give you the best service, it helps a lot if you read the following submission advice.

- Please be sure to include any relevant details of your system. 'I can't get X to work' doesn't really mean anything to us if we don't know things like what version of X you are trying to run, what hardware you are running on.
- Be specific about your problem. Things like 'it doesn't work' or 'I get an error' aren't all that helpful. In what way does something not work? What were you expecting to happen? What does the error message actually say?
- Please remember that the people who write this magazine are NOT the authors or developers of Linux, any particular package or distro. Sometimes the people responsible for software have more information available on websites etc. Try reading the documentation!

We will try and answer all questions. If we don't answer yours specifically, you'll probably find we've answered one just like it. We can't really give personal replies to all your questions.

WRITE TO US AT:
 Linux Format, Future Publishing, 30 Monmouth Street, Bath BA1 2BW or email: lxf.answers@futurenet.co.uk

ANSWERS



Urchin web stats

Q I don't know if what I am asking is possible or makes any sense but you guys seem to be good at interesting work-arounds. I have installed a web stats program called *Urchin*. It has its own web server (which runs on port 9999 in my case). I want to run *Urchin* on the same physical server as all my websites. It is possible to run *Urchin* program on a different port number but one of my bigger customers only allows access to port 80 through their corporate proxy server.

I also have the same problem with *Webmin* which runs its own server, but I want to run it on port 80 with a different domain name. Is there any way around this apart from having to assign another IP address? I cannot configure them for port 80 or my main *Apache* daemon will not start.

Name and address withheld

A Your main *Apache* server can act as a gateway to other web servers. This is generally

referred to as a reverse proxy. To configure this, you need to configure a virtual host in *Apache* that will do the forwarding.

If your *Urchin* server is running on port 9999, the section in your `httpd.conf` file should look like the following:

```
<VirtualHost *>
    DocumentRoot /var/www/html
    ServerName
webstats.yourdomain.com
    ErrorLog
logs/webstats.yourdomain.com-
error_log
    CustomLog
logs/webstats.yourdomain.com-
access_log common
    ProxyPass / http://localhost:9999/
</VirtualHost>
```

Also make sure the proxy module is being loaded earlier in the configuration file, for example:

```
LoadModule proxy_module
modules/libproxy.so
AddModule mod_proxy.c
```

Now make sure a DNS entry exists for `webstats.yourdomain.com` that points to the IP address of your server.

ABOUT RACKSPACE MANAGED HOSTING

RackSpace Managed Hosting™ provides complete managed hosting solutions with Fanatical Support to small to medium-sized customers as well as large enterprises. All customer platforms include state-of-the-art data centres, customised servers, burstable connectivity, 99.999% uptime SLA, a dedicated account manager, instant

emergency response and access to live expert technicians 24x7 for support of all hardware and core software. Founded in 1998 with locations in San Antonio, Texas and London, the company manages over 7,000 servers for customers in more than 80 countries. For more information, visit www.rackspace.co.uk and see this month's *Linux Pro*.



FREE UP MEMORY

Q I have been running a Linux server for only a short time and am very new to the whole way of administrating Unix-like OSes. I recently found the 'top' command and noticed I only have a few megabytes of memory free out of my 512MB of memory. I think this is unusual as my system is not really being used much yet as we are still learning how to use it. I've also verified the information with *free*. I've included the relevant information from *top* and *free*:

	total	used	free	shared	buffers	cached
<i>free:</i>						
Mem	514692K	494104K	20588K	OK	22056K	343572K
<i>top:</i>						
Mem	514692K	494508K	20184K	OK	22060K	OK
Swap	1004052K	OK	1004052K	OK	OK	343756K

From the LXF forums

A *top* shows you two different types of memory usage. Memory usage, and Swap Usage (virtual memory). The way that Linux works is that the Memory usage includes what's called 'cached' or Cache Files. In Linux, the Cached memory is used to speed up your system (as with many other operating systems). It does this by storing files that are used/accessed regularly in memory. This number is constantly being tuned to your system's usage. As your system requires more memory, the cached variable will tend to become slightly lower; in some cases as it frees up memory in order for applications/daemons/processes etc to use more. This is normally done in the background so you wouldn't really notice it. By looking at your free output, you can get an idea of what's using up your memory. To find out what's in your cache you can list open files by using the tool *lsof*.

You will now be able to access your web stats program by simply connecting to this address:

<http://webstats.yourdomain.com>

Idle CPU

Q I have recently been experiencing very slow responses from my server.

However when I go into *top*, I find that the CPU is almost idle (10-15% utilised). I have an Athlon 2000 with 256MB RAM serving webpages from *PHPNuke*. Why would this be so slow?

From the LXF forums

A With only the limited information you've supplied about your system, I can't really give you a firm answer; however, I would start by looking into your virtual memory (swap) usage. You will probably find that your system is swapping large amounts of memory in order to serve content. *PHPNuke*, being a database-based application, will use a lot more memory than a standard static html service. To verify this yourself you can look at using *vmstat*. This tool will give you understandable information on all of your virtual memory.

Database backup

Q I would like to do a backup of MySQL database. Is dump a good tool to use and how would I go about doing this. Is it better to use tar?

I want to backup my whole system (running Red Hat 9.0) to an extra hard drive.

From the LXF forums

A The problem with backing up databases as files is that the content of the file may change during the actual backup procedure. The ideal thing to do would be to export the database to a text file and back up this 'snapshot' instead. If you're running the backup job on a regular basis using *cron* you can add this facility into the *cron* script.

The choice between *tar* and *dump* is a personal choice; both are very capable but have different strengths that can be used in different circumstances. *dump* is designed to duplicate entire file systems. This may not be ideal depending what files you want to backup and depending on your file system layout. *tar* allows backing up of individual files or directories, so may be better if you just want to back up your databases.

WIN A SHARP ZAURUS PDA FROM RACKSPACE MANAGED HOSTING

www.rackspace.co.uk

Every month, the best question related to Systems Administration that a LXF reader sends in wins a Sharp Zaurus SL-5500 running Linux.

With industry support from the likes of IBM, SAP, Oracle, and Sybase, the Zaurus features a QWERTY keyboard behind a sliding cover. The 65,536 colour screen is 240x320 resolution. An Intel® StrongARM® 206MHz processor and 64MB memory makes the Zaurus ideal for downloading software or playing MP3 and MPEG1. The docking station that fits into its USB connector, allowing users to synch with their desktop, schedules, address

book and ToDo lists. The email app supports POP3 or IMAP email servers. Write and send emails and browse the Internet with a compatible modem card, or using the infra-red port and a mobile phone; or any replied, forwarded, or composed emails are transferred to your outbox the next time you synch. SD and CF slots allow upgrades and addition of software.

Email all your sysadmin questions and problems to Hans Huberland: sysadminqa@rackspace.co.uk Every published query on these system admin pages wins a Rackspace T-shirt.

It's only 74mm wide by 138mm long and a svelte 18mm thick – so the Sharp Zaurus really is Linux in your pocket!



★ Star Question – PDA winner!

This issue's lucky winner is *Nils* – your new Zaurus PDA will be with you shortly!

Backup solutions

Q I work for a small company for which I installed a 1TB file server replacing a 240GB server. At the moment the server is using 236GB of the 1TB, but to back this up as it stands now will take about 10 hours and 10 DDS4 Tapes – quite a big job, I'm sure you'll agree. The tape drive is a HP DDS4, which has a max of 20GB with 2:1 compression per tape; it's not an autoloader so each tape has to be put in manually. I googled around but I find tape solutions to be very expensive about 4-5 thousand Euros for a 300GB tape backup system. It will be my head if anything goes wrong (can't blame it on Bill Gates this time!).

With the previous machine a full backup would take about 8-9 hours at about 80% disk space utilisation. That would be every month so it wasn't too bad with incremental backups throughout the weeks.

You are probably saying to yourself "why don't you just backup old stuff and take it off?" As the company is an animation company, jobs can span weeks or even months, so the information needs to be accessible at a 'double-click'. I was toying with the idea of creating another file server and using that as a mirror of the main one so if the main server went down. It would only take a few minutes to be back online, but then if there were a fire etc, this server would be gone too.

Being a hosting company Rackspace would have to deal with this everyday, what in your opinion would be the best (cheap) solution for a Tape System?

Nils, via email

There is a very handy utility called `mysqlhotcopy` for performing backups on databases; the syntax as follows:

```
/usr/bin/mysqlhotcopy dbname  
/backupdirectory
```

In this example, the database name being backed up is `dbname` and it will be backed up to `/backupdirectory/dbname`.

Device file deleted

Q I accidentally deleted the device file for my RAID array of 2 disks striped (/dev/md0). The array is still there according to the kernel boot messages, but obviously I cannot mount it. How can I make the device file again, or is there any

A Due to the unusual storage space requirement of your company and the budget restraints you've described, there is no ideal solution. Backing up to tape sounds great but backup hardware is very expensive, more so in the case of very high capacity devices. If a 300GB single tape device is too expensive, then there is very little else you could do to pursue the tape backup strategy. Also bear in mind that restoring such quantities of data from tape can take even longer and this would be made worse if using incremental backups.

A second server provides very fast backups and ease of recovery should the primary server become unavailable. It would have the operating system and data already in place and this eliminates the need to reload the OS and then try to recover from slow tape media. Using a second server means that if there is a fire and the entire building burns down, you are left with nothing. Moving the second server offsite means you would need connectivity to the server and fast lines are expensive, much more expensive than a very-high end tape backup solution.

Another option would be to buy a large IDE or SCSI hard drive and add it into your server. You could mirror the data onto this disk fairly quickly and then remove the drive and keep it in a safe place. This solution is cheap but remember that hard drives are not the most reliable media, especially when you're taking them home each weekend. Depending on your server hardware you may need to reboot your server to add or remove a hard drive each time and this may not be acceptable either.

As you can tell each of these options has its own merits and pitfalls in the categories of price, reliability, ease of use and speed. A good solution will not be cheap but your company needs to decide how much more it would cost if the data were lost for whatever reason.

other way to fix this? These drives contain projects that I need for college, so any help you can suggest is greatly appreciated.

From the LXF forums

A The command `mknod /dev/md0 b 9 0` issued as the root user should recreate the entry.

This will recreate the device `/dev/md0` as a block device with a major number of `9` and a minor number of `0` (the major and minor numbers specify the type of hardware you are trying to use). You should check that the ownership is `root:root` and you will probably want the permissions to be `rw-rw-rw-`



Don't throw away your old PC – run a lightweight distro on it and/or turn it into a network client.

« Legacy Hardware

I have an old laptop that I'd like to install Linux on. It's a 486/66 with 24MB memory and a 500MB hard drive, so I'm thinking an old version of *Slackware* (I have install CDs for Slackware 3.3, Red Hat 4.2 and Debian 1.3.1, any of which should run OK on this machine) – anyone got any better suggestions?

But now for the interesting bit... The machine has no CD drive, and the floppy drive is kaput. It does however have a working PCMCIA Ethernet card, so I should be able to install across a network, if I can get it to boot an installer.

It currently has Windows 95 on it (upgraded with floppies from DOS/3.1 before the floppy drive died). Anybody got any bright ideas as to how I can get it to boot and start an install across a network without using the floppy drive?

Nigel, From the LXF forums
Most Linux distros support network installs, although it depends if you can actually boot the laptop off the network, or if you trust Windows 95 to work well enough to boot a Linux kernel using **loadlin**. Of course, the big problem is that if your Linux installation doesn't boot properly after the install, then you're going to be a little stuck.

Doing a network install using BOOTP, DHCP and NFS is possibly doable, if the NIC supports network booting. Unfortunately, most PCMCIA cards don't, as the PCMCIA subsystem doesn't kick in until the OS gets going.

inetd or not inetd?

On the man pages there is usually a line telling you the command line

options to run the daemon (*ftpd*, *httpd*, *sshd* etc) as a standalone process or you can start it by using *inetd* instead. What is the best way, standalone or *inetd*? My progress through trying out daemons was:

- 1 *ftpd* and *telnet* through *inetd*
- 2 Stopped using *ftpd* and *telnet* through *inetd* and now running *sshd* as a standalone process since I can do both with SSH (on a Windows box when I am away from my machine)

3 Planning to add *httpd* next
From the LXF forums

The concern when running processes through *inetd* is that you have one process forked per connection, which may be a bad thing. In the case of SSH, each connection would require it to build its keys all over again, which may take 10-20 seconds on a slower system. If you run it standalone, it'll build the keys once when the box boots up and each connection to it will use those keys. *Apache* is very similar, in that running it through *inetd* would require it to parse out its configs and open file descriptors to the log files for each request or connection. This is a major performance issue, so running *Apache* through *inetd* is usually not recommended.

As *telnetd* doesn't do much of anything when you start it, it's perfectly acceptable to run it through *inetd*. FTP is also similar to *telnetd*, although as it starts a separate connection for the actual data transfer, the effects of lots of connections via *inetd* are not noticeable. However, some servers such as *proftpd* work better standalone, as they are a little more involved than basic FTP servers. **LXF**

TURNING JAPANESE

Pictographical fonts in X

Switching between

I am currently using SUSE 8.1 Professional, and I am trying to set up my computer to work bilingually with English and Japanese. Would you please tell me whether I will be able to input Japanese text to a program while running an English desktop/OS; and if I can, how do I go about it?

From what little information I have been able to find on the Internet, it is necessary to have a conversion engine/daemon running, and be able to start X as a child process of an IME. With the a little poking during the install on SuSE, I have been able to get Canna running in the background as it should, but I do not know how to make sure that the IME is running correctly and find a way to let me switch between Japanese and English. I have even done the Japanese installation, but it still seems to leave my input as English characters-only. It is my understanding that a special keyboard is not required, and the IME should handle the process of converting English character text to Japanese characters.

Mark Biesiada, via email

Bilingual

I have asked this on a few forums but have not had any satisfactory replies. I am trying to wean my Japanese wife from Windows. She is keen to use Linux but she needs to be able to use English for her work as well as Japanese. She gets a mixture of Japanese and English email, which she needs to respond to accordingly.

I have tried Red Hat, Mandrake and SUSE (all latest versions) and set up KDE and GNOME to use Japanese as the main language so that you can shift between them via the keyboard, but this does not work. I have even installed them as Japanese versions but they still type in English only.

With Windows, you have IME running at the bottom of the screen and you are one click away from changing the input from

English to Japanese. Any suggestions as to how I can make this happen?

Paul Taylor, via email

While localisation in many Linux programs can often be highly sophisticated and complete, using Japanese – or any pictographical fonts – in X in conjunction with English is not the easiest thing to do, although it can be done. It is documented in some depth at www.ugcs.caltech.edu/~diffusor/xim/. Both a conversion daemon and an input engine is required to be installed:

```
# apt-get install canna canna-utils
libcanna1g
# apt-get install kinput2-common
kinput2-canna
$ export
XMODIFIERS="@im=kinput2"
$ kinput2 -canna &
```

The following also needs to be added to *.Xresources* or *Xdefaults*, depending on the distribution:

```
*.inputMethod: kinput2
*.preeditType: OverTheSpot
Kinput2*useOverrideShellForMode:
true
Kinput2*conversionStartKeys:
Shift<Key>space
```

```
! for Kterm
KTerm*VT100*KanjiMode: euc
```

```
! for rxvt
Rxvt*multichar_encoding: eucj
```

```
! for Gvim
Vim*international: True
Vim*inputMethod: kinput2
Vim*preeditType: OnTheSpot
```

If you're in X, run `% xrdb -merge .Xresources` to have your changes take effect.

To start inputting text in Japanese, **shift-space** will allow any Japanese-aware X application to handle the appropriate characters. Lots of tips, information and some example configurations are provided on the site listed above. While it is Debian-specific, it is fairly straightforward to obtain the appropriate packages for a different distribution and perform the same function without having to resort to Debian if that is not your choice of distro.

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).



November 2003

Product code:
LXFB0046(cd)
LXFD0046(dvd)

DVD HIGHLIGHTS:
RUNT, RPM Analyser,
Linux Easy Access
Keyboard, VCool

CDs HIGHLIGHTS:
Gentoo Linux, Seapine
Surround SCM FULL
VERSION, Nano, Mass
Rotate, BZFlag,
Jahshaka, Kroupware,
ZoneMinder, BatchGIMP,
ConfluxGroupware

MAGAZINE FEATURING:
The new 2.6 kernel, Java IDEs
Roundup, PHP 5 preview,
Linux keyboards, Aethera, The
Grid, OOo 1.1, Armari Pro3D,
Regular Expressions tutorial,
New-look SUSE, ZoneMinder



October 2003

Product code:
LXFB0045(cd)
LXFD0045(dvd)

DVD HIGHLIGHTS:
Topologlinux, JRE SDK,
Wifiscanner, Knokisync,
Squirrel SQL Client

CDs HIGHLIGHTS:
FreeBSD4.8, K3b,
Kognition, Beneath A
Steel Sky, ScummVM,
Thunderbird & Firebird,
Java2, Java Runtime
Environment, Scribus
1.0, LM Sensors

MAGAZINE FEATURING:
Speed up Linux, Red Hat
Package Management part 2,
Server School tutorial,
OpenGroupware.org review,
Which webserver, What On
Earth is BSD? Stop spam



September 2003

Product code:
LXFB0044(cd)
LXFD0044(dvd)

DVD HIGHLIGHTS:
JAMD Linux, LimeWire,
MoviX, KOffice 1.3,
Scribus, LBreakout2,
Mozilla 1.4, Opera 7.11

CDs HIGHLIGHTS:
MULTIMEDIA SPECIAL
GNUSound, TuxPaint,
Cinelerra, Realsoft 3D,
AlsaPlayer, TreePad,
Ecamegapedal, Grip,
Kino, Xine, TiMidity++

MAGAZINE FEATURING:
40 pages of tutorials, SuSE
Enterprise Desktop 1 review,
Text Editors Roundup, Linux
Sound feature, Seapine SCM,
JBuilder9, Navaho TeamCAT,
Neoware Thin Client review



August 2003

Product code:
LXFB0043(cd)
LXFD0043(dvd)

DVD HIGHLIGHTS:
GnuCash, 3dDesktop,
Slackware Live CD, LDP,
Astaro Security Linux,
SharpROM v3, Cinelerra,
ScummVM, TrustixSL

CDs HIGHLIGHTS:
Complete Networking
Toolkit, SpamAssassin,
AGT, GuardDog, Jwall,
Nagios, America's Army,
Samba (stable & dev)

MAGAZINE FEATURING:
Networking feature, VMware
review, CD-burner roundup,
Evesham ELinux desktop PC,
Opteron servers, Linux training,
TextMaker 2002, NetVault 7,
GnuCash, 8 beginners' pages



July 2003

Product code:
LXFB0042(cd)
LXFD0042(dvd)

DVD HIGHLIGHTS:
Opera 7.1 beta, GAIM,
ABCDE, Pingus, KDE-
Sounds/Aqua Fusion,
LDP, Alsaplayer, KOffice,
RUTE, Seclude, Freevix,
DamnSmallLinux

CDs HIGHLIGHTS:
Morphix, GuardDog,
Nvidia XDrivers, GCC
3.3, TORCS, Linux
Cookbook, SETI

MAGAZINE FEATURING:
All about X, MySQL review,
Sharq server review, WINE
explained, Steganography,
Distributed Computing, Red
Hat Enterprise AS, Choosing
a distro, GCC 3.3 review



June 2003

Product code:
LXFB0041(cd)
LXFD0041(dvd)

DVD HIGHLIGHTS:
Slackware 9.0, Mutt,
OpenOffice.org, Xine,
TakeAJoint, Peanut
Linux, FCKeditor, KDE
3.1.1a, mirror of Linux
Documentation Project

CDs HIGHLIGHTS:
Mandrake 9.1, AC3D,
KDEradio, The GIMP –
including 400 plug-ins
from The GIMP Registry

MAGAZINE FEATURING:
AMD's 64-bit Opteron CPU,
Linux Laptop Megatest, The
Hurd, JCodeBox review, Build
your own RPMs, Red Hat 9
vs SuSE8.2 vs Mandrake 9.1,
OpenAuth Linux PC review

Other back issues also available, please call for further information, giving the issue number and/or date required

To order, call the back issues hotline on **0870 4448645**

or **+44 1458 271178** for overseas orders

Remember to quote the issue code shown above and
have your credit card/switch card details ready.

Prices (inc post and packing): CD – UK£5.99,
US & Eur £6.99, rest of world £7.99
DVD – UK £6.99, US&Eur £7.99, rest of world £8.99.

LINUX FORMAT

www.linuxformat.co.uk

User Groups

LUGs worldwide are full of members keen to help with your problems, discuss ideas, and generally natter about all things Linux. You can find lots more information online at: www.lug.org.uk

1 HAMPSHIRE

URL www.hants.lug.org.uk
Contact Hugo Mills

2 BRISTOL & BATH

URL www.bristol.lug.org.uk

3 SCOTTISH

URL www.scottish.lug.org.uk

4 OXFORD

URL www.oxford.lug.org.uk
Contact Alasdair G Kergon

5 KENT

URL www.kent.lug.org.uk
Contact Kevin Groves

6 BRIGHTON

URL www.brighton.lug.org.uk
Contact Johnathan Swan

7 WORCESTERSHIRE

URL www.worcs.lug.org.uk

8 NORTHANTS

URL www.northants.lug.org.uk
Contact Kevin Taylor

9 ANGLIAN

URL www.anglian.lug.org.uk
Contact Martyn Drake

10 MILTON KEYNES

URL www.mk.lug.org.uk
Contact Denny De La Haye

11 SCUNTHORPE & DONCASTER

URL www.scundog.org
Contact Shaun Holt – shaun@scundog.org

12 MORAY

URL www.moray.lug.org.uk
Contact Stewart Watson

13 WEST WALES

URL www.westwales.lug.org.uk
Contact Dan Field

14 WOLVES

URL www.wolveslug.org.uk
Contact Jono Bacon

15 PETERBOROUGH

URL www.peterboro.lug.org.uk
Contact Steve Gallagher

16 EDINBURGH

URL www.edinburgh.lug.org.uk
Contact Alistair Murray

17 TYNESIDE

URL www.tyneside.lug.org.uk
Contact Brian Ronald

18 LEICESTER

URL www.leicester.lug.org.uk
Contact Clive Jones

19 GREATER LONDON

URL <http://glug.linux.co.uk/>
Contact John Southern

20 SURREY

URL www.surrey.lug.org.uk
Contact Jay Bennie

21 CAMBRIDGE

URL www.cam-lug.org.uk

22 DEVON & CORNWALL

URL www.dclug.org.uk
Contact Simon Waters

23 FALKIRK

URL www.falkirk.lug.org.uk

24 MANCHESTER

URL www.manlug.mcc.ac.uk
Contact John Heaton, Owen Le Blanc

25 HERTFORDSHIRE

URL www.herts.lug.org.uk
Contact Nicolas Pike

26 WEST YORKSHIRE

URL www.wylug.lug.org.uk
Contact Jim Jackson

27 SHEFFIELD

URL www.shefug.co.uk
Contact Richard Ibbotson

28 STAFFORDSHIRE

URL www.staffslug.org.uk

29 NORTH EAST

URL www.shofar.uklinux.net/NELUG

30 LONDON

URL www.lonix.org.uk

31 BERKSHIRE & THAMES VALLEY

URL www.sclug.org.uk

32 LIVERPOOL OPENSOURCE

URL http://linux.liv.ac.uk/_liv_linux_ug/
Contact Simon Hood

33 DEAL AMIGA CLUB

Email superhighwayman@hotmail.com
Contact John Worthington

34 CHESTERFIELD

Email spirelug@yahoo.co.uk
Contact Robin Needham

35 SOUTH DERBYSHIRE

URL www.sderby.lug.org.uk
Contact Dominic Knight

36 BELFAST (BLUG)

URL www.belfastlinux.cx
Email russell@belfastlinux.org

37 WILTSHIRE

URL www.wiltshire.lug.org.uk
Contact Jason Rudgard

38 SOUTH LONDON

URL www.sl.lug.org.uk
Email edo@perceptiondm.com

39 CHESHIRE

URL www.sc.lug.org.uk
Contact Anthony Prime – enquiry@sc.lug.org.uk

40 NORTH WALES

URL www.northwales.lug.org.uk
Contact Andy Hutchings A-Wing deltaone@virgin.net

41 MIDLANDS

URL <http://midlandslug.port5.com/>
Contact Pete Thompson

42 CUMBRIA

URL www.cumbria.lug.org.uk
Contact Jamie Dainton

43 DORSET

URL www.dorset.lug.org.uk
Contact John Robinson

44 SHROPSHIRE

URL www.shropshire.lug.org.uk
Email shropshire@lug.org.uk

45 SOUTH WEST

URL www.southwest.lug.org.uk
Email southwest@lug.org.uk

46 SOUTH WALES

URL www.swlug.org.uk

47 NORTH LONDON

URL www.kemputing.net/lug/anlug-aims.html
Email jason@voyagercomputers.co.uk

48 MALVERN

URL www.malvern.lug.org.uk
Contact Greg Wright

49 HUDDERSFIELD

URL www.hud.lug.org.uk
Contact Dave Naylor – knocker@caramboo.com

50 NOTTINGHAM

URL www.nottingham.lug.org.uk

51 ST ALBANS & LUTON

URL www.lust.lug.org.uk
Contact Michael Culverhouse – mike@easily.co.uk

52 WREXHAM

Contact Paul Kersey-Smith
Email paul@pkls.fsnet.co.uk

53 PRESTON & LANCS

URL www.preston.lug.org.uk
Contact Phil Robinson

54 DERRY

URL www.derry.lug.org.uk

55 ISLE OF WIGHT

URL www.iow.lug.org.uk
Contact David Groom – info@iow.lug.org.uk

56 SCARBOROUGH

URL www.scarborough.lug.org.uk

57 BLACKBURN

Email matt@consultmatt.co.uk

58 YORK

URL www.york.lug.org.uk

59 LINCS

URL www.lincs.lug.org.uk

**60 HULL**URL www.hull.lug.org.uk**61 WALTON-ON-THAMES**Contact William Mutch
Email rael@freeuk.com**62 GLOUCS & COTSWOLDS**URL www.gloucs.lug.org.uk**63 WEST OF SCOTLAND**URL www.wos.lug.org.uk**64 SOUTH STAFFORDSHIRE**URL www.staffs.lug.org.uk**65 MANSFIELD**URL www.mansfield.lug.org.uk**66 BORDERS**URL www.linux.bordernet.co.uk**67 BIRMINGHAM**URL www.sb.lug.org.uk**68 COVENTRY**Email info@coventry.lug.org.uk**69 NEWARK & LINCOLN**URL www.newlinc.lug.org.uk**70 BEDFORDSHIRE**URL www.beds.lug.org.uk**71 LINCOLN**URL www.lincoln.lug.org.uk**72 LOUGHBOROUGH**URL www.loughborough.lug.org.uk**73 EXETER UNIVERSITY**URL www.euslug.lug.org.ukEmail N.J.Murison@exeter.ac.uk**74 SUNDERLAND**Email thomas.croucher@sunderland.ac.uk**75 EAST YORKSHIRE**Email sharkonline@whatemail.com**76 CLEVELAND OPEN SOURCE GROUP**Email openlug@digitalmedia.co.uk**77 BEVERLEY**Email vladimir_lukyanov@hotmail.com**78 DUNDEE & TAYSIDE**URL www.dundee.lug.org.uk**79 SUSSEX**URL <http://sussex.lug.org.uk/>**80 WIGAN & ST HELENS**Email paulf.johnson@ukonline.co.uk**81 BRIXTON**URL www.communitytechnology.org.uk/~linuxhome**82 ST.ANDREWS, FIFE**URL www.standrews.lug.org.ukEmail stuart@nx14.com**83 NUNEATON**URL www.nuneaton.lug.org.uk**84 ISLE OF MAN**URL www.iom.lug.org.ukEmail helix@manx.net**85 AYLESBURY**URL www.aylesbury.lug.org.ukEmail drbond@educational-computing.co.uk**86 LANCASHIRE**URL www.lancasterlug.org.uk**87 EAST LONDON**URL www.eastlondon.lug.org.uk

Contact Jonathan Spriggs

88 ORMSKIRKEmail rob@northwestlinux.co.uk**89 HEREFORD**URL www.hereford.lug.org.uk/Email rbjh@good-news.fsnet.co.uk**90 EAST HERTS**Email madtom1999@yahoo.com**91 SWINDON**Email nick.trueman@ntlworld.com**92 MENAI**URL www.menai.lug.org.uk**93 ABERDEEN**URL www.aberdeen.lug.org.uk**94 SHETLAND**URL www.shetland.lug.org.ukEmail c_s_s_butler@yahoo.com**95 GLASTONBURY**URL www.glastonbury.lug.org.uk

Contact Steve Leonard-Clarke

96 SOUTHEND-ON-SEAURL www.sos.lug.org.uk

Contact Derek Shaw

97 ORPINGTONURL www.orpington.lug.org.uk

Contact Barry Schofield

YOUNG LINUXURL www.young.lug.org.uk**SCHOOLS**URL www.schools.lug.org.uk

LUG OF THE MONTH

UK Unix User Group (UKUUG)

Call for participation:

UKUUG will hold its next LISA/Winter conference in Bournemouth, Dorset on 25 and 26 February 2004. The event's theme is High-Availability and Resilience; and will take the form of a series of presentations, each followed by a discussion on issues raised. To this end, papers are invited from interested parties on the general themes and related topics.

The UKUUG wishes to encourage discussion on all aspects of systems

and their administration; and is especially interested in papers covering theory and practice, high-availability, performance, network management, novel solutions to practical problems, integration, interoperability and security (including the business, legal and moral issues). As always, UKUUG wishes to encourage work-in-progress presentations, and student project posters; proposals for these should be submitted in the same way as for full papers. You do not have to be a

member of UKUUG to submit a paper, and submissions from speakers outside the UK are welcome. Potential authors may request further information at winterconf@ukuug.org; initial abstracts should be sent to the same address. Abstracts should be accompanied by a short biography, and ideally should be about 250-500 words. Final papers should normally last 30-40 minutes. If you need more time for your presentation, please tell us when you submit your extract.

SIGNIFICANT DATES:

Closing date for abstracts:

11 November 2003

Authors notified:

18 November 2003

Programme published:

End November 2003

Final papers due:

13 January 2004

See www.ukuug.org for more, and membership details etc should be sent to office@ukuug.org.

Worldwide Linux User Groups

Free Software users across the globe

Africa**EGYPT**

URL www.linux-egypt.org

GAUTENG, SOUTH AFRICA

URL www.glug.org.za

Email glugmin@revolution.org.za

THE LORD'S ABODE, JO'BURG, SA

Email Andrew Gargan avrin17@iname.com

Australia**ADELAIDE**

URL www.linuxsa.org.au

Email mtippet@anu.edu.au

ALICE SPRINGS

URL www.aslug.org.au

MELBOURNE, VICTORIA

URL www.luv.asn.au

Contact luv-committee@luv.asn.au

PERTH

URL <http://plug.linux.org.au/>

SYDNEY

URL www.slug.org.au

Europe**COSTA DEL SOL (English speaking)**

URL www.fuengirola.lug.org.uk

DENMARK

AIssund www.alslug.dk

Esbjerg www.eslug.dk

Fyns www.flug.dk

Midt-og Vestjylland www.mvjlug.dk

Nordjylland www.njlug.dk

Skåne Sjælland www.sslug.dk

Trekantsområdet www.tlug.dk

Vest-fyn www.haarby-net.dk/vflug

Århus www.aalug.dk

EIRE

URL www.linux.ie

Email root@linux.ie

URL www.dilu.org

Contact glossary@dilu.org

MILUG (Longford)

URL <http://midlands.linux.ie>

Contact midlands@linux.ie

Middle East**ISRAEL**

URL www.iglu.org.il/IGLU/

Contact webmaster@iglu.org.il

PALESTINE

URL www.lugps.org

Email isam@planet.edu

Asia**HONG KONG (multilingual)**

URL www.linux.org.hk

SINGAPORE – SLUG

URL www.lugs.org.sg

SRI LANKA

URL www.lklug.pdn.ac.lk

MYANMAR (formerly BURMA)

URL www.myanmarlug.org

Email afityde@balug.org

PAKISTAN

URL www.linuxpakistan.net

Email tux@cslug.org

HYDERABAD, SINDH, INDUS VALLEY

URL www.geocities.com/slug_pk/

KASHMIR

Coming soon!

China

BEIJING (GB encoding, but mostly written in Chinese)

URL <http://mud.263.net.cn/~linux>

CHINESE LINUX USER GROUP

URL www.linux.org.cn

NANJING

URL <http://jllib.jlonline.com/njlug>

India**LINUX INDIA**

URL <http://linux-india.org>

ALIGARH LUG

URL <http://linux.amupost.com>

BOMBAY

URL www.ilug-bom.org.in

CHANDIGARH

URL www.geocities.com/vipinb

CHENNAI AND MADRAS

URL www.chennaiug.org/

CYBERABAD (CLUG)

URL <http://seeknew.freesevers.com/clug/>

DELHI

URL www.linux-delhi.org

KOLKATA

URL www.ilug-cal.org

MADURI

URL <http://linuxmadurai.tripod.com>

NORTHERN INDIA LINUX

URL <http://groups.yahoo.com/group/lug-northindia>



Spreading the word

“How can Linux make *my* business more profit?” Jono Bacon explores the pitfalls of advocating Linux use in enterprise.

Business are an incredibly important target market for Linux. In the next few issues we will look at different types of businesses and how they can be focussed on, but we will generally identify the small-to-medium enterprise (SME) market. This market typically includes small startups and mid-sized operations that do not have bottomless pockets.

Selling Linux as a platform to these businesses involves two core concepts that are strengths within the Linux sphere: cost and stability. Many companies are looking for a complete system to run their business but generally do not want to pay through the nose for it. As such, if someone goes up to a business and suggests a solution that is not only suitable for their needs but is low cost, they will not believe their luck, especially when they are then made aware of the stability benefits.

Sadly, the people responsible for making decisions on IT expenditure often have little or no awareness of the computing market outside the Microsoft Windows world, so you will need to tread

very carefully in your advocacy – you're not likely to convert anyone to the Linux way of doing things if your initial presentation or suggestions make them look like an ill-informed idiot!

The key in this area is knowing what the business needs to achieve. This may be a variety of functions, and the platform may need to be able to talk to other systems such as Windows/Mac/Unix machines. Always remember to be honest and do not pledge that Linux can do something it can't. When you are clear of the benefits, you then need to back these benefits up with statistical info that shows that Linux is not only a popular and capable system, but one that others have successfully made the jump to already. Businesses need peace of mind in their systems, and this statistical info is particularly important. You will also have to carefully explain “who to blame” when things go wrong, as Open Source/Free software is still widely misunderstood.

Next month we will look more at businesses; comments and suggestions to spreadingtheword@jonobacon.org **LXF**

Linux User Group organisers

If you're not listed here, or we have your details wrong, please contact us at: **LUGS!, Linux Format, 30 Monmouth Street, Bath, BA1 2BW** or email your details to: linuxformat@futurenet.co.uk

Essential disc info

Read this important information before you use your *Linux Format* coverdisc – CD or DVD. We've collated some helpful info to help you get the most from these jewels of data!

FINDING THE ESSENTIALS

Missing something?

As many of the programs on our discs are the very latest releases, they are often built on the very latest libraries and may depend on other packages your current Linux setup does not contain. We try to provide you with as many of these important supporting files and libraries as possible, though obviously we don't have space to include absolutely everything.

In many cases, the latest libraries and

other packages you might need will be included in the "essentials" folder on the disc, so if you are missing dependencies, this is the first place to look.

Package formats

Wherever possible, we try to include as many different types of package for an installation as possible, whether that be distribution specific RPMs, debs or whatever. Please bear in mind that we can only do this where space permits and when the packages are available.

We will, apart from exceptional or legally restricted situations, include the source files for any package, so that you can build it yourself.

Documentation

These pages provide helpful information on how to install and use some of the packages on the CD. Please note that many of the applications come with their own documentation, and there are additional notes and files in the relevant directories.

WHAT ARE ALL THESE FILES?

If you are new to Linux, you may find the profusion of different files and extensions confusing. As we try to give as many packages as possible for compatibility, there will often be two or three files in a directory covering different types of Linux, different architectures and usually source and binary versions – so which do you install? They can be identified by their filenames, and usually just by the file extensions.

Someap-1.0.1.i386.rpm – This is probably a binary rpm, designed to run on x86 systems.

Someap-1.0.1.i386.deb – The same, but a debian package.

Someap-1.0.1.tar.gz – This is usually source code.

Someap-1.0.1.tgz – Same as the above, tgz is abbreviated form of tar.gz

Someap-1.0.1.tar.bz2 – Same, but uses bzip2 compression instead of zip

Someap-1.0.1.src.rpm – This is also source code, but supplied as an rpm to make it easier to install

Someap-1.0.1.i386.RH7.RPM – A binary, x86 RPM designed specifically for Red Hat Linux

Someap-1.0.1.ppc.Suse7.rpm – A binary RPM designed specifically for SuSE7.x PPC Linux.

Someap-devel-1.0.1.i386.rpm – A development version.

INSTALLING FROM TARBALLS

A tar ball is a two stage archive. First the files are archived into a single file with tar and then compressed with Gzip or Bzip2. To unpack, cd to the directory you want to unpack it, usually your home directory and type one of the following two lines:

```
tar xzvf /mnt/cdrom/Desktop/progname/progname-2.1.0.tgz
```

```
tar xvf --bzip2 /mnt/cdrom/Desktop/progname/progname-2.1.0.tar.bz2
```

Use the first for Gzipped files, those ending in .tar.gz or .tgz, and the second for Bzipped files, ending in .tar.bz2 or .tbz2. Naturally, you change the paths to suit the location and name of the archive. and replace /mnt/cdrom with whatever is applicable to your system (eg /cdrom). This normally unpacks the archive into a directory of the same name, enter that directory with:

```
cd progname-2.1.0
```

To compile and install the software, type the following three commands:

```
./configure
```

```
make
```

```
su -c "make install"
```

The last line will prompt you for the root password, as this stage must be run as root. If you are already logged in as root, just type **make install**. This will give you a default installation. If you want to change any aspect of the install, type

./configure --help to see the options available. For example, you are usually able to change the default location with the PREFIX argument. When you have finished installing, you may remove the source files with:

```
cd ..
```

```
rm -fr progname-2.1.0
```

You should also log out as root, before you do anything you may later regret.

DEFECTIVE CDs

In the unlikely event of your disc being defective please email our support team (support@futurenet.co.uk) for further assistance. If you would prefer to talk to a member of our reader support team please call **01225 822 743**.

CREATING INSTALL CDS WITH CDRECORD

The quickest way to burn an ISO image to CD is with *cdrecord*. You need to be root to do this. First find the address of your CD-writer with

```
cdrecord --scanbus
```

This will show the devices connected to your system. The SCSI address of each device is the three numbers in the leftmost column, say 0,3,0. Now you can burn a CD with

```
cdrecord dev=0,3,0 -v /path/to/image.iso
```

You can simplify the command by saving some default settings in

`/etc/default/cdrecord`. Add a line for each CD writer on your system (usually one) like this

```
Plextor= 0,3,0 12 16M
```

The first item is a label, after the SCSI address you put the speed and the buffer size to use. You can now replace the SCSI address in the command line with the label, but it gets even easier if you add

```
CDR_DEVICE=Plextor
```

Now you can burn an ISO image to disc with

```
cdrecord -v/path/to/image.iso
```

If you really don't want to use the command line, *gcombust* will do the job for you. Start it as root, select the 'Burn' tab and the 'ISO 9660 Image' gadget near the top of the window. Put the path to the image file in the gadget and press 'Combust!'

Now put on the kettle while the CD is created for you.

Other OS?

You don't have to use Linux to burn the ISO to a disc. All Linux-specific bits are already built into the image file. Programs like *cdrecord* simply dump it to the disk. If you don't have a CD-writer, find someone who has one, and a DVD drive, and use the CD burning software on their computer. It can be Windows, Mac OS, AmigaOS whatever.

No CD burner?

What if you have no CD writer? Do you know someone else with one? You don't have to use Linux to burn the CDs, any operating system that can run a CD-writer will do the job (see above).

With some distributions it is also possible to mount the images and do a network install, or even a local install from another disk partition. The methods often vary between distributions, so check on the distro vendors website for more information. [LXF](#)

Coverdisc



Neil Bothwick used *MailScanner* this month to stop his system from drowning under a tsunami of Sven virus-originated emails – for more info on *SpamAssassin*, see page 94!

There's not so much space available for discussing the contents of the DVD this month, with so much else to write about, but here's a quick run-down of a couple of the highlights, that we've managed to squeeze in around the edges of Mandrake 9.2!

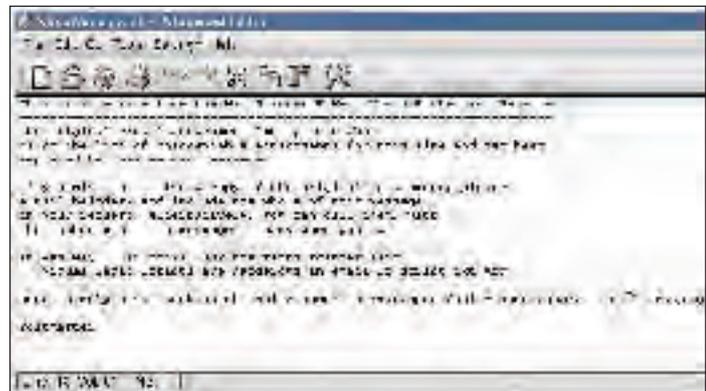
SERVER MAILSCANNER

Anyone who tells you that Linux is immune to viruses is plain wrong. While the risk of being infected is low, we are still affected by those running more vulnerable systems. I've been receiving up to 5MB of mail per day from

machines infected with Sven. This is bad enough if you only run Linux, especially if you have a dial-up connection to the Internet. If you have one or more Windows machines on your network, the position becomes far more serious, as they could become infected.

MailScanner is an email spam and virus detection package. Unlike some other such programs, *MailScanner* is not targeted at one specific mail server or virus scanner. It works with a range of different servers and various commercial and free virus scanners. It also links up with *SpamAssassin* for extremely effective spam detection.

In addition to detecting and filtering unwelcome mails, *MailScanner* is able to modify the content of mails, removing images from pornographic spam to keep obscene content away from your network and removing code known to exploit vulnerabilities of some mailers. It is also able to keep your installed virus scanners up to date, by contacting their update servers every hour. Once an hour may seem extreme, but recent viruses have spread so quickly that it is vital to keep your database as up-to-date as



MailScanner will protect your network from email-borne viruses and spam. There is a Webmin configuration module included too.

possible. The RPM packages are contained in .tar.gz archives. Each archive contains a number of RPMs and an install script.

DESKTOP XFCE

There is a mass of publicity every time a new version of the KDE or GNOME desktop environments is released, and you have already read about the inclusion of GNOME 2.4 on the CDs. These are both powerful and flexible

desktops, perhaps a little too powerful for some people's tastes. There are also some far more lightweight, and minimal, desktops around, like *Fluxbox*. Somewhere in between falls *XFCE*. It is not as resource hungry as the two main heavyweight contenders, while still retaining a lot of functionality and flexibility. There are two reasons you may want to try a more lightweight desktop. It could be that you are using a system with less power than today's multi-GHz monsters. Or it could be that you have plenty of power but want to save as much of it for running applications rather than displaying their interfaces.

XFCE 3 did a job job of providing the compromise between power and efficiency, the new version 4.0 looks set to continue and improve on its heritage. This is not just a window manager, but a full desktop environment along the lines of GNOME and KDE.

The tarball unpacks to a number of separate source tarballs, which should be installed using the normal **./configure && make && make install**. The correct order is given in the install.html file. To install to Mandrake 9.1 or Red Hat 9, cd to the relevant directory and type, as root `rpm -Uhv *.rpm`
The Mandrake RPMs will probably

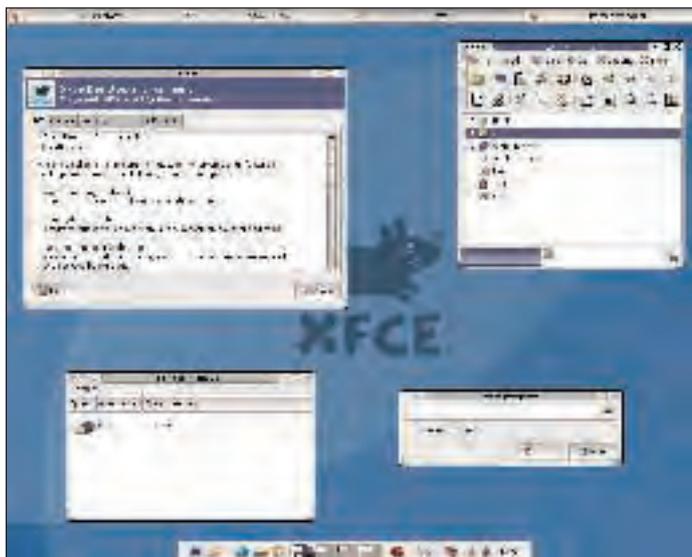
On the DVD



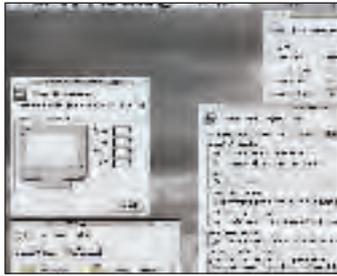
Wherever you see this logo it means there's related stuff on the DVD

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.



XFCE 4.0 is far more than a window manager. Here are some of the desktop configuration programs.



XFCE 4.0 is fully internationalised – this is a desktop in Tamil!

work with 9.2, but this hasn't been tested yet. To install for any other RPM distribution, you should build your own RPMs from the packages in the SRPMS directory.

HELP LDP

Every month we carry a complete mirror of the Linux Documentation Project on the DVD. This extremely useful online resource carries an immense number of HOWTOs, tutorials, man pages, articles and other information. It has been a while since it was highlighted in these pages. In itself that is reason enough to remind you of its presence and usefulness. However, there is another reason for mentioning it now. The Linux Documentation Project is celebrating its first decade – Happy Birthday!

It was ten years ago that Matt Walsh wrote the Installation HOWTO, the first ever Linux HOWTO. From this, the project has grown into one of the largest collaborative projects on the Internet. The project is currently home to nearly 500 documents, which are the work of around 450 writers and translators and appear in fourteen

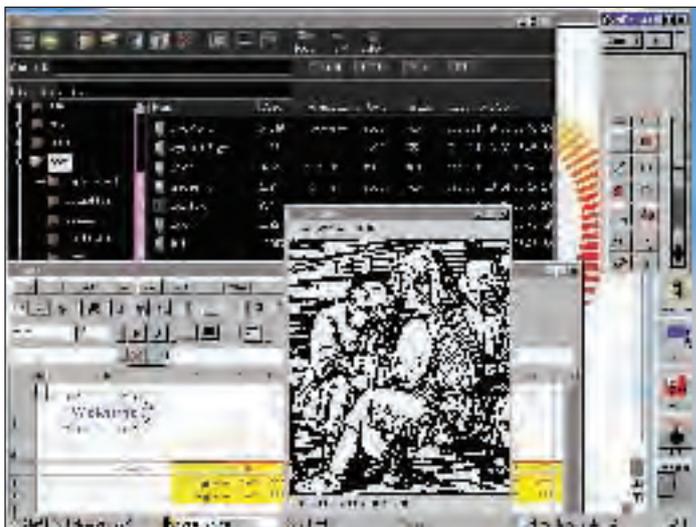
languages. If you have a question, the chances are that someone else has already asked the question and written down the answer. It's not unusual for the results of a Google search on a Linux question to have more than half of its front page results on the www.tdlp.org website or its mirrors.

DISTROS PUPPYLINUX

Following the recent surge of interest in live CD distros, Linux distributions that boot and run from a full-sized CD, we are now seeing an increasing number of mini-distros. Klaus Knopper created a great deal of interest in live distros with his excellent Knoppix, and now people are taking advantage of the wide and cheap availability of more portable storage media to produce distros that fall somewhere between a full CD and a floppy disk.

Puppy Linux is one such distro. At less than 30MB, it can boot from a mini CD-R, Compact Flash card or USB device. Despite the minimal size, Puppy Linux contains an incredible amount of software, including a paint program, DTP, web browsers and spreadsheets.

The PuppyLinux directory on the DVD contains an ISO image to create a CD. Just burn it in the usual way, as described on the Essential Disc Info page. If you want to use a different kind of device, such as a USB flash disk, you will need to copy some files to the disk before setting it up as bootable with syslinux. Full information on how to do this is contained in flash-puppy.htm, but it refers to files that do not appear to be on the DVD, vmlinuz



Boot a remarkably full-featured Linux desktop from a 32MB flash disk. Puppy Linux comes with full instructions to make your distro-on-a-keyring.



Slashdot was one of very many sites to report on, and complain about, Verisign's infamous DNS wildcard. Protect your local DNS service against this and future such abuses, with *dnsmasq*.

and image.gz. These files are contained in the ISO filesystem, it would have been a waste of space to include separate copies as well, so you need to mount the ISO image with `mount -o loop -t iso9660 /mnt/cdrom/Distros/PuppyLinux/cd-puppy.iso /mnt/point`

Then you can copy the files from /mnt/point to your USB device. This is an early version, not yet 1.0, so hardware support is still limited. The only way to tell if it will work on your computer is to try it. I tried it on three computers and it worked on two of them. The one that failed used a VIA EPIA M-series motherboard with onboard everything.

GRAPHICS FLPHOTO

There are dozens of programs that will display images or create photo albums. We're not saying this one is the best, but it is good. Its author describes it as "a simple image management, viewing, and printing program", which pretty well sums up what it does. It provides many of the basic functions that are needed when dealing with collections of images. These are the most often needed basic functions, collected together in one program. It doesn't try to compete with heavyweights like *The GIMP* for image manipulation, but it does provide useful functions like rotating, scaling, brightness adjustment and red eye removal.

There are functions to organise images into albums, and photos can be imported directly from digital cameras using *gphoto2*, which almost certainly came with your distribution. Albums can be viewed on screen, displayed as a slideshow or printed. Albums can also be exported to web

pages, so the whole World can see pictures of your new baby/pet/car/boat.

SYSTEM DNSMASQ

A couple of unpleasant things have happened to the Internet in the last few weeks. One was the Sven virus, mentioned earlier, and similar malware. The other was the way Verisign kludged the master DNS records so that any request for an invalid, or mistyped, .com or .net domain resolved to their SiteFinder service, rather than returning the error code many programs expect. BIND, a popular Internet DNS server, was patched to refuse this address, in order to restore some semblance of normality, which is fine if that's the software your ISP uses.

Many people use *dnsmasq* to provide domain name service for a local network, BIND is rather extravagant for a small LAN, and it has been on previous cover discs. *Dnsmasq* allows one machine to provide DNS information for the local network from its /etc/hosts file, while forwarding all other requests to a server on the Internet. This latest version includes a fix for Verisign's attempt to hijack the DNS service, by interpreting the address of their SiteFinder site, the address that is returned when no such domain exists, as an error. This has been done in a configurable way, so any address that may be used in such a way in future can be easily blocked.

To enable this feature, install the new version and add a line to /etc/dnsmasq.conf that specifies the address to block. For Verisign's site, the line to add is

```
bogus-nxdomain=64.94.110.11
```

To block future DNS abuses, just change the IP number in this line. [LXF](http://www.linuxformat.co.uk)

LINUX PRO

FROM THE MAKERS OF LINUX FOR

DECEMBER 2003



PLUS

David Valentine

IBM's EMEA Linux supremo, outlines the new server strategy and development plans

Linux: desktop saviour?

Find out why Wyse is pushing Linux solutions to this niche but rapidly growing market

Rethink Storage

Xinit's Sharqstore 3TB NAS – more space... lower costs

Server hardening

In order to fix the problems, you need to know what's wrong!

YOUR WEB WORRIES SORTED!

Do you meet the features of Apache, but with the 24-hour support and the manageability of a commercial solution? It's time you checked out Covalent...

PRACTICAL LINUX SOLUTIONS FOR I.T. PROFESSIONALS

Welcome

Twenty pages of real-world Linux for IT professionals

Summer, we can reveal, is officially over. So why is there a guy lapping up rays as he reclines on some pleasant pasture adorning the cover? This is the ideal that server administrators hanker for – no worries, no-effort maintenance. And for the more picky of you, with remote admin tools, the server admin could be enjoying what promises to be another hot summer in Cape Town. But how do we get to this ideal? By choosing smart services and smart strategies. That's one of the things *Linux Pro* is here for!

Apache is one of the leading lights of Open Source. As well as being notably brilliant, it is also probably the most widespread single Open Source project used in business. But unless you have the in-house expertise and support to use it effectively, it may not be the best solution. That's where Covalent steps in, with the nearest thing to a commercial version of the world's best web server. Alternatively, why run your own servers when there are dozens of helpful hosting companies to do it for you? Our series of Profiles continues this issue with a look at Rackspace.

Server admin isn't an easy job at the best of times, but increasingly much time and effort is taken up trying to prevent the more malicious elements of society from doing things they shouldn't with your boxes, or worse, trying to clean up in the aftermath of such an event. Keeping on top of vulnerabilities and patches is a tedious but necessary business. But you don't need to be alone! Turn the tools of the cracker against them. *Nessus* is a scanner, but one designed with the aim of providing security analysis and help for hardening your servers. Being well-prepared is often the difference between inconvenience and disaster!

As ever, we welcome your feedback on *Linux Pro* – if there are any subjects you would like to see covered here, mail me!

Nick Veitch Editor
nick.veitch@futurenet.co.uk



“This is the ideal server administrators hanker for – no worries, low-effort maintenance. But how do we get to that place?”

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

Group Publisher

Dave Taylor

Photography

Powerstock

ADVERTISING

Senior Sales Executive

Andrew Hudson

andrew.hudson@futurenet.co.uk

(01225) 442244 Ext 2441

Copyright No part of this publication may be reproduced without written permission from our publisher. All letters & contributions to *Linux Format* are submitted and accepted on the basis of non-exclusive worldwide licence to publish or license others to do so unless otherwise agreed in advance in writing. *Linux Format* recognises all copyrights in this issue. Where possible, we have acknowledged the copyright holder. Contact us if we haven't credited your copyright and we will always correct any oversight. We cannot be held responsible for mistakes or misprints.

LINUX is a trademark of Linus Torvalds, GNU/Linux is abbreviated to Linux throughout for brevity. All other trademarks are the property of their respective owners.

LinuxPro is a part of



magazine

Published by



Media with passion

30 Monmouth Street, Bath BA1 2BW
Tel +44(0) 1225 442244

CONTENTS

COVALENT

Web-serving without worries **p4**

DAVID VALENTINE

IBM's Linux strategy revealed **p8**

THIN CLIENTS

Wyse's vision for the desktop **p10**

XINIT SHARQSTORE

Affordable NAS storage **p14**

HARDENING WITH NESSUS

Scanning your own servers **p16**

HOSTING PROFILE

Spotlight on Rackspace **p18**



COVER FEATURE **COVALENT**



COVER FEATURE
Your web worries
SORTED

If you're serious about web serving, you need to find a company capable of giving you the software and support necessary to fit all your needs, not just some.
PAUL HUDSON discovers why you need to talk to Covalent...



When it comes to web servers close to our hearts, most of us would mention the venerable *Apache* first, and it's not surprising – *Apache* is one of the most capable, mature, and speedy web servers available. However, the natural downside to it being a free software project is that it doesn't come with any guaranteed support as part of the package, which leaves many people uncertain about deploying it. While it's true that *Apache* does have mailing lists, and does have IRC channels where you can go and ask your questions, you're not guaranteed an answer, and, even if you will eventually get an answer to your question, there's no guarantee that it will work as expected.

So, what do you do when you want the power of *Apache* but with the addition of reliable technical support? Back in January 1998, some of the key developers of *Apache* were thinking just the same thing, and the end result was Covalent – a corporate face to *Apache*, designed to add

extra security, reliability, and manageability to *Apache*. Nowadays, Covalent employs such greats as Jim Jagielski (charter member of the *Apache* project and co-developer of *NCSA Server*), Doug MacEachern (author of *mod_perl*), Ryan Bloom (a leading developer of *Apache 2*), and more. These people still actively contribute massive amounts of work back to *Apache*, making sure that Covalent gives back to *Apache* as much (if not more) than it takes.

With *Apache* deployed on 65% of web servers around the world (source: Netcraft), it's clear that *Apache* is very popular. However, because Covalent's software reports itself back as vanilla *Apache*, what *isn't* clear is quite how widely deployed Covalent's software has become: 175 of the Fortune 500 customers run Covalent. Perhaps the key selling point for Covalent is that it provides 24x7 support from its team of engineers – a feature that most users will find invaluable when faced with one of *Apache*'s more-or-less inexplicable error messages. 

STRATEGIES AND SOLUTIONS

We spoke to Jim Zemlin, the VP of Marketing at Covalent, about products, plans, and the technology driving it all.

LINUX PRO: In your opinion, what are the primary reasons why a company should choose Covalent to formulate its web solution?

JIM ZEMLIN: No other company offers a management application that covers the entire Open Source stack from *Apache* to *Tomcat* to MySQL to Linux to *JBoss*. No other company offers an application-centric approach to managing web infrastructure that will dramatically improve application availability and response times.

LXP: How important was the release of *Apache 2* to you as a company?

JZ: It was important to us because our customers had been demanding some of the features in *Apache 2.0* for quite some time, particularly some of the performance-related improvements around multi-threading. Having said that, we still do a strong business supporting both *Apache 2.0* and *Apache 1.3*.

LXP: Where does .NET and J2EE fit into Covalent's strategy?

JZ: We consider ourselves neutral. Covalent's management software supports both J2EE and .NET applications.

LXP: To what extent is Covalent moving into other areas?

JZ: Covalent has recently launched our Covalent Application Manager product, which leverages our experience managing the web server and applies it to the entire application stack. New customers of this product are significantly reducing the time it takes for them to resolve problems and greatly improving service levels around response time and availability across the board.

“We are involved in the Apache Software Foundation, working on improving the core server and giving the code back to the ASF.”

JIM ZEMLIN, COVALENT



LXP: To what degree are you involved with the development of Web standards?

JZ: As you know the Web – and *Apache* itself – is reliant on standards. We support all major standards in our products. As a management vendor, we must do that. Many of our engineers are involved in the Apache Software Foundation, which works in setting standards.

LXP: So you give back a lot of code to the *Apache* movement?

JZ: We have employees that work on improving the core *Apache* server and give back all that code to the ASF. The only code we keep are changes we make to our proprietary management tools which do not require modifications in the *Apache* code at all because they are written using the *Apache* modular API.

LXP: How much demand is Web application management in?

JZ: Over that past several years, Covalent has witnessed the incredible explosion of web

application technology. Those deployments are now online and need management. Just as the need for storage management followed the build-out of large storage capacity, web application management is following an unprecedented build out of web apps.

LXP: Where do you see web server/application management and related paradigms going in the next few years?

JZ: The market for management solutions for Web applications only stands to increase. In fact, many analysts proclaim that only 25% of Web applications in development have been deployed. That's a huge market. For the future, we see a rise in predictive management technologies. Instead of waiting for your Web app to break, wouldn't it make sense to solve potential problems before they occurred? That's why management solutions that integrate analysis with monitoring have the advantage. It's not enough to monitor applications, you need to optimise their performance and their costs for greatest business impact.

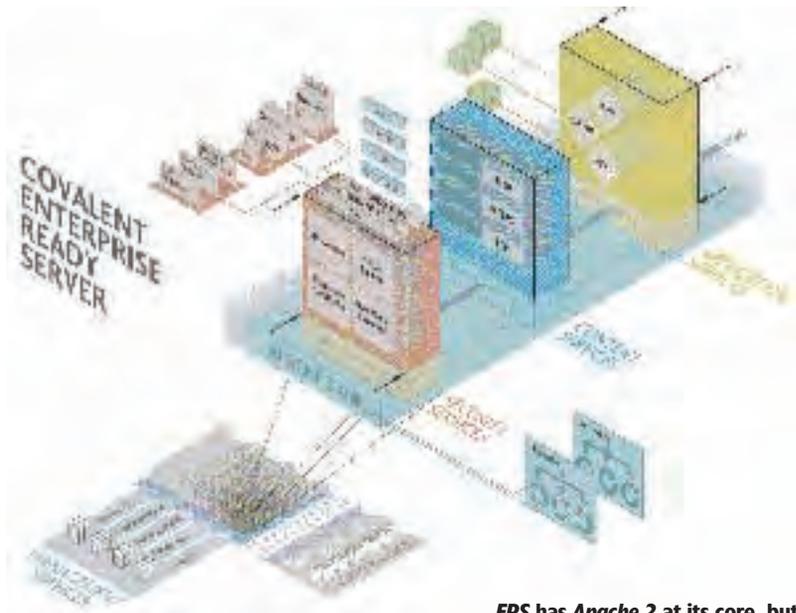
COVER FEATURE **COVALENT**

◀◀ The product line

Although the core of Covalent's product revolves around *Apache*, the company has added a lot of value on top of it beyond just offering technical support. The flagship product is *Covalent Application Manager (CAM)*, which is a comprehensive management solution combining all core operational activities for managing web applications into one solution. The goal of *CAM* is to shift web infrastructure management from component-specific to business-oriented application management via its proprietary application model, thus fulfilling Covalent's apparent mission statement of "manage applications, not servers". *CAM*'s application model, is designed to allow companies to manage their web infrastructure and thereby improve service levels, reduce costs and drive enhanced communication between IT and the business.

For companies that don't need to control applications from a central server, Covalent also produces *Enterprise-Ready Server (ERS)*, which is designed to tailor *Apache*'s featureset towards enterprise customers, and, of course, comes with Covalent's full-time support. While not as feature-rich as *CAM*, *ERS* does still have a lot to offer, particularly when it comes to security – it has dozens of security-related tweaks, from the separation of authentication and authorisation to the incorporation of the high-speed *RSA* security libraries. Covalent also produces *Enterprise FTP Server*, also based on *Apache*, which can work as a standalone solution if desired, or can be configured to work as an add-on to *ERS*.

The final product in the Covalent suite is *Covalent Fast Start Server (FSS)*, which is designed to be a fairly basic



ERS has *Apache 2* at its core, but adds extras such as single sign-on and *RSA* libraries.

Apache server that walks you through installation and configuration of *Apache* with a friendly GUI – you can choose which modules you want installed, what security you want, and so on. Customers who find their business outgrowing *FSS* can easily upgrade to *Enterprise-Ready Server*, making *FSS* a powerful little package that will serve the needs of most SMEs.

That said, *CAM* adds so much to the equation that its popularity is unsurprising once you get into the arena of

COVALENT IN PRACTICE

To find out how well Covalent's plans translate to real-world differences, we spoke to Michael Freed, the Web Project Leader at the Iowa Foundation for Medical Care, where Covalent's software is currently handling its infrastructure.

LINUX PRO: What made the Iowa Foundation for Medical Care consider Covalent?

MICHAEL FREED: Cost and support were the two biggest issues. We had developed prototypes using the open source versions of *Apache* and *Tomcat* but our Government clients were concerned about these two issues.

LXP: How easy was it to get the right solution?

MF: Since our prototype application was developed to the Open Source versions of *Apache* and *Tomcat*, the transition was nearly seamless. We purchased multiple licenses of *Covalent Enterprise-Ready Server* and have been very happy with both the performance of the product itself and Covalent's support as well.

LXP: What did you find most difficult? What did you find most beneficial?

MF: Our migration was pretty simple and rather painless. The only problems we've experienced were with the early versions of the Management System.

If you made direct changes to any of the *Apache* config files, the management system (which provides a GUI interface to the config files) would not recognise those changes and that resulted in some misconfigurations. This is all fixed now.

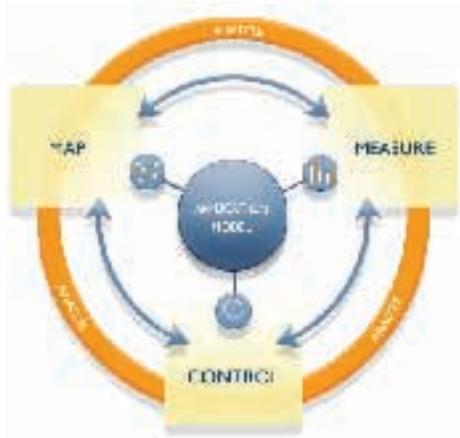
LXP: What was the biggest advantage?

MF: We utilise Covalent to host multiple government web sites ranging from an award winning secure site to a content management system site. Prior to migrating to Covalent we used *Oracle 9iAS Release1* which was much more expensive and riddled with problems. Running on *9iAS* we were bombarded by *JServe*-related performance issues and a significant number of security issues. Oracle support's answer was "wait and buy the next version."

Since migrating to Covalent, we have not experienced one minute of security- or attack-related downtime. Our record with *9iAS* was not as shining. At that point our decision was clear, many of the market leaders were building their application servers around *Apache*. We wanted that solid and proven foundation with someone who'd truly back that up and not charge an enormous amount of money for less-than-adequate support.

The biggest advantage of migrating to Covalent thus far has been its support. The company has been very responsive and have worked closely with us when necessary. Given the large amount of money many companies charge for poor support these days, it's not hard to calculate significant ROI on this move.

“Since migrating to Covalent, we have not experienced one minute of security-related or attack-related downtime.”



This diagram makes *Application Manager* look simple, but behind the scenes a lot of work goes on to ensure heterogenous compatibility is maintained.

larger corporates – it allows sysadmins to easily map applications and supporting resources, to measure application service levels and resolve problems quickly, and analyse application data for reporting, forecasting and planning purposes – all features not yet seen in competing products.

Who uses Covalent?

Covalent's client list reads like a *Who's Who* of successful companies, and includes GE, US Bank, Hewitt, AMD, Deutsche Bank, AT&T, and GlaxoSmithKline. The list could go on and on – NASA, HP, Johnson & Johnson, Lucent, and thousands of others all weighed up the service that Covalent provides and decided it was the best move for them. With such a clear enterprise lead, it's no surprise that so many companies and research groups openly acknowledge Covalent's position, such as the director of Intel's Solutions Development Group, Michael Pope, who said "Companies such as Covalent are leading the way in building sophisticated web server software".

With the web application market now on the verge of major growth, the keyword seems to be "flexibility" – more companies than ever are now considering deploying web applications, and Covalent has positioned itself as the definitive answer for web-based IT operations in order to be ready for the surge. A key part of this positioning is partnership, and Covalent is certainly no slacker here – they recently announced a partnership with BEA Systems to ensure that the BEA's *WebLogic Platform* works seamlessly with the *Covalent Application Manager*. This latest agreement adds to the raft of existing partners, including IBM, HP, Sun, Intel, Verisign, and Computer Associates.

In order to educate customers on how to customise their Covalent products for their own purposes, Covalent also offers training for its products, and also generic *Apache 2*-based training, which is unique in the marketplace. This training can either be online, for customers who want to pace their learning, or onsite, where Covalent's engineers come to your office to provide a customised training solution.

Covalent vs Apache

In a business with veteran sysadmins, what are the compelling business reasons for choosing to pay money for Covalent as opposed to installing *Apache* for free? While the knowledge that you have 24x7 support available whenever you need it might be nice to know, it won't

feature high on your list if you already have a trained team. Instead, the key advantages will probably be the stronger security model, the integrated application management tools, and the comprehensive technology support for just about every popular web-based product out there.

Although Covalent's explanatory diagrams often do more to confuse than to help, the reality is that their software really does increase the performance, reliability, stability, and security of the *Apache* server, partly by adding support for a number of security techniques, including adding full support for hardware encryption, but partly because they have their finger so close to the pulse of the *Apache* team and also to the stock market – they know open source and they know the enterprise.

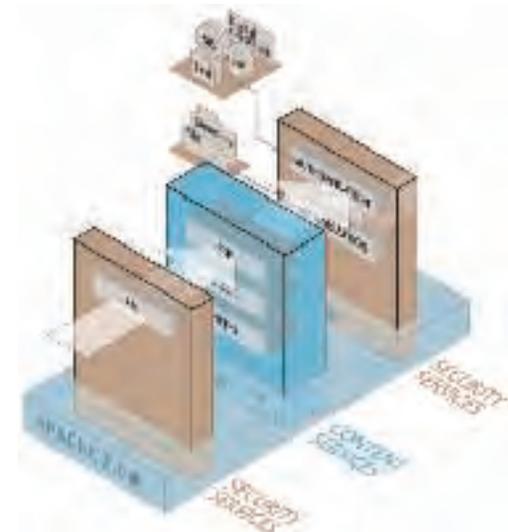
While some people might argue that choosing Covalent over *Apache* doesn't help Open Source, the reality is again quite different – Covalent employs several key *Apache* developers, and all of the modifications they make to the *Apache* core are contributed back to the community.

Future plans

Covalent has firmly committed itself to enhancing *Covalent Application Manager* so that it continues to grow with the demands of its customers. As web applications continue to edge into the spotlight, new standards are being designed and approved to help make management easier. To keep up with all this change, Covalent is nearly always in the process of refining and updating their offerings – soon to be added, for example, is support for AIX, MySQL, and *ColdFusion* from directly within *CAM*. In order to help third-parties develop their own resources for use within *CAM*, Covalent is in the process of opening up their application definition, with the eventual goal of making it an open standard.

Despite having experienced widespread take-up already, *CAM* is almost certainly going to see even larger growth in the next year to eighteen months as more companies make the commitment to web services. Right now, and probably for the foreseeable future, Covalent simply has no competition in this arena, which leaves the company with a hot product in what is sure to soon be a very hot topic. ■

Covalent's FTP Server is also based up on *Apache*, and can easily be integrated with *ERS*.



Linux Technology Center & Government Solutions Center: IBM's enterprise strategy for Linux

LINUX PRO: What was IBM's Unix history before Linux?

DAVID VALENTINE: IBM began shipping its Unix products in 1990 when it launched the RS/6000 product range, running the AIX operating system. During the 1990s, this server technology rapidly gained market share becoming one of the premiere Unix offerings in the marketplace. In 1999 the product was integrated into a single IBM server brand, and is now designated the IBM eServer pSeries, where 'p' stands for 'performance'. These servers are running in thousands of customer sites around the world, often running mission critical applications for customers, and the IBM-designed POWER microprocessors they run have become a standard for 64-bit computing in the computer industry, with similar processors appearing in games modules, and engine management systems.

LXP: What motivated IBM's interest in Linux?

DV: In the mid 1990s, we made a commitment to deliver open standards based products and solutions. Linux, coming from the Open Source community, and supporting open standards was a natural choice. In discussions with our customers, we found that many large organisations, such as banks, were starting to use Linux as a cost-effective, reliable alternative to proprietary Unix platforms for web serving, file and print, email, and other infrastructure applications.

The customers were attracted to the potential cost savings and performance of deploying Linux for some apps, but wanted the high levels of reliability support they had come to expect from IBM. As customers began to demand open standards for software integration, Linux became a logical choice for the implementation of more sophisticated solutions, and IBM uniquely could offer real choice, able to offer the very best platform for the specific type of work, running on mainframe servers, or Unix servers, or Intel or AMD based servers, running either 32- or 64-bit microprocessors.

LXP: How many people does IBM have working on Open Source projects?

DV: IBM has had a large team of developers working on Linux and related Open Source projects for many years now. We call this the IBM Linux Technology Center (LTC) and today it numbers over 250 people in various locations around the world including the UK. They work on a long list of projects that are typically aimed at making Linux suitable for enterprise deployment. Recent areas of development

Linux Pro spoke to DAVID VALENTINE, the Linux Sales Manager for EMEA, about IBM's Unix history, how it has been applying that to Linux, and where IBM thinks Open Source software is going...

include high availability and symmetrical multi processing scalability. IBM engineers are leading many open source projects, with the large number of IBM sourced research papers submitted each year for the Ottawa Symposium, a clear demonstration of the contribution IBM is continuing to make to the Open Source community.

LXP: How does Linux fit into IBM's server strategy?

DV: IBM is committed to delivering the very best technology and solutions for our customers, and as we pioneer breakthroughs in server capabilities, price performance and server design, we enable new workloads for customers, that were not economically viable before, or just not technically possible before. Many of these new applications are Linux-based, exploiting the eServer technology advantage for competitive gain.

In the field of supercomputing, industries and organisations such as government, academia, life sciences, automotive, oil exploration, aerospace, digital media creation, and business intelligence have benefited from breakthroughs in AMD and Intel microprocessor design, with IBM developing complete Linux-based cluster solutions to radically transform customer computing.

Cars and aeroplanes can now be designed in cyberspace, automotive crash tests can now be simulated to dramatically speed up development and reduce costs, and the latest movie releases can economically have thousands of computer-generated scenes in them.

Another area where we are seeing breakthrough technology is in the area of virtualisation, where it is now

IBM was a platinum sponsor and keynote speaker at London's Linux Expo UK 2003.



possible, using our high end IBM eServer zSeries mainframe, running the z/OS operating system, to host hundreds of instances of z/Linux virtual servers on a single computer, overcoming the existing scalability issues associated with Linux, with additional servers able to be provisioned in less than 30 minutes, and able to be reclaimed when not required almost instantaneously, so capacity can be adapted to changing business needs.

IBM foresees Linux spreading throughout the business world at an unprecedented rate, in all industries, including telecoms, petroleum, geophysical sciences, financial services, government, automotive, computer aided engineering and design, as the economics of Linux are becoming more and more attractive for some of our customers.

LXP: Do you think Linux deserves more of a place on the desktop, or is it just a server OS?

DV: IBM's Linux effort is focused on Linux & OSS solutions that bring business benefits to our customers. Recently Linux has been receiving a lot of consideration as a desktop OS, especially within some governments, such as in Germany. There is also a trend towards Linux for specific function desktop systems, such as engineering workstations, or for bank tellers. IBM will continue to work with our customers in all areas where Linux & OSS could be of benefit.

LXP: How important is kernel 2.6 to IBM's plans?

DV: The 2.6 kernel will be a breakthrough not just for IBM but for the entire Linux community, with new levels of functionality expected to further enhance Linux adoption for many of our customers.

LXP: How committed are IBM to the future of Linux?

DV: IBM is fully committed to the future of Linux. More than 5,000 IBM employees work on Linux in research, services, development, porting centres and sales & marketing. Our entire eServer range, all five of our middleware brands (Lotus, DB2, Tivoli, Websphere, Rational) support Linux today. We continue to release more products with Linux, and contribute to the Open Source community.

There are now more than 6,300 IBM Linux customer engagements worldwide, allowing customers to reduce their computing costs with solutions ranging from web serving to some of the largest supercomputers doing seismic processing, financial calculations and genomic research. Worldwide, more than 75 IBM government customers – including agencies in France, Spain, UK, Australia, Mexico, the United States and Japan – have embraced Linux to save costs, consolidate workloads, increase efficiency and enact e-government transformation.

Over 80 solutions have been announced by Solution Providers for Linux on iSeries. There are now more than 50,000 Windows and Intel developers actively creating Linux-based apps that run on IBM software, including WebSphere, DB2, Lotus and Tivoli, to build applications that run on Linux. These developers have created more than 6,500 Linux-based apps for IBM software. Some 4,700 Business Partners support Linux-enabled IBM software. Two-thirds of these new Linux-based apps are being created by corporate developers, signalling that more and

“A large team is working on Linux and related Open Source projects – IBM's LTC (Linux Technology Center) numbers over 250 people.”

David Valentine, IBM Linux Sales



more businesses are making the move to Linux. One third are being created by ISVs. IBM is shipping over 65 software products on Linux across its IBM DB2, WebSphere, Lotus and Tivoli software families.

IBM is committed to using Linux inside IBM with more than 3,500 servers running Linux. Mission-critical apps that run Linux include the IBM website, support of IBM's new \$2.5Billion 300mm chip manufacturing facility, and apps supporting more than 300,000 IBM employees worldwide. In addition, IBM hosts websites for many of its customers on Linux, including Wimbledon, the US Open, the French Open and many other sporting events. IBM has opened a variety of Linux-based centres around the world that are dedicated to helping customers, Business Partners and developers move to an open, standards-based approach to computing. From IBM's Linux Center of Competence on Wall Street to the Government Solutions Center in Washington DC, IBM gives developers a place to test new Linux-based solutions.

LXP: To what extent do your servers take advantage of new hardware, such as Opteron?

DV: IBM's server strategy for Linux has always been to offer customers choice through responding to their demands. The IBM eServer 325 which incorporates the AMD Opteron processor can either be bought outright by the customer or they have the option to access them through IBM's deep computing on demand facility in Poughkeepsie, NY, paying for processing power based on the required capacity and duration of use.

Designed to run either Linux or Windows, the IBM eServer 325 provides high performance computing customers with increased performance and seamless migration from 32- to 64-bit server technology. The IBM eServer 325 is also part of IBM's Cluster 1350, providing a complete cluster solution including a broad portfolio of clustering software in common with our cluster offerings on POWER and AIX/Æ.

Technologies such as the IBM eServer 325 powered by the Opteron processor offers strong performance and extended memory addressability while ensuring backward compatibility that preserves customers existing 32-bit software investments. IBM is the first major IT vendor to take advantage of this latest technology and deliver it to our customers. ■

Wyse up to thin clients

With the entry of Wyse into the Linux thin client market, *Linux Pro* re-appraises this niche but important sphere.

Since the 1980s, the desktop PC revolutionised the corporate world. Soon, almost every worker in every area of business found a need for a desktop computer, and the progressively cheaper desktop machine fulfilled the need admirably.

The dumb terminal connected to a powerful central server still had its place, but it was standalone desktops that found their way into business. That is still pretty much the world we live in today, only more so. The spread of the Internet as an underlying technology for modern business has meant more people than ever have found a need for a computer at work – it's hard to think of many office occupations that don't require a computer these days.

But in many ways, the standalone PC is a poor choice for the corporate environment. As any sysadmin will no doubt know, it is a solution which puts too much power in the hands of the people they trust the least – users. The IT landscape of the office is not a harmonious, homogenous technological idyll, but a federation of semi-autonomous machines operated by people with vastly different levels of computer literacy and training.

It isn't surprising then that server-oriented corporate computing is on the rise once more. Thin clients are being



Some users think they're being short-changed when they see the small hardware footprint!

PROS AND CONS OF THIN CLIENTS

PROS

- + Usually Cheaper to implement
- + Reduced IT management overhead
- + More effective use of resources
- + Scalable
- + Various remote admin options

CONS

- Not suitable for some high performance computing tasks
- Limited usability of external devices
- High server/network demand
- Can introduce a single point of failure – eg if the network/server goes down

sold in larger volumes all over Europe and North America. These solutions – reworking of the terminal concept – are proving popular because of the great advantages, primarily in terms of administration. As the hardware itself has few (if any) moving parts, and provides an identical hardware base that can be simply locked down, suddenly there is a lot less need for engineers to go out on site to fix problems, often not caused by problems with the IT systems themselves, but by user tampering.

IT departments have never been busier than they are right now, so any time-saving advance is going to have both a cost and efficiency benefit to any organisation that is looking to adopt the thin client way of doing things.

Thin advantage

With a thin client solution, though there may need to be a consequent reinvestment in terms of networking infrastructure, the long-term benefits are clear in a whole range of scenarios.

In today's network-driven world, the majority of computer tasks are carried out with a very small subset of apps – an email client and web browser with Java support fulfils most people's needs, perhaps with the addition of a small number of office apps such as a wordprocessor and spreadsheet.

Many of the advantages of thin clients play well to the strengths of Linux itself. Obviously, there is the underlying stability of the OS itself, advantageous in any circumstances. Another prominent consideration is cost. Licensing costs for proprietary operating systems drive up the cost of each unit, and may have further cost implications in upgrades, patches etc. With Linux there are no such worries.

But it's the often-overlooked aspect of the OS that have a particular impact on this market. As an Open Source OS, Linux itself is very flexible and configurable. From the kernel itself to the supporting OS tools, everything can be selected and tuned to match the hardware perfectly, resulting in a more efficient OS that fits into a smaller footprint (and consumes less system memory), again with cost savings for the hardware.

On the application front, Linux clients can make use of the long established Citrix server software to run Windows applications that may be needed, but as most office functionality is based around web-based applications, this may not even be necessary. On the flip-side, legacy server-based apps that require X11 or terminal modes are well catered for with Linux, as one might expect.

Basics

The basic requirements of a thin client unit are fairly minimal. A

WYSE WORDS

Linux Pro speaks with Stephen Yeo, European Marketing Director, Wyse Technology, to discuss the future of thin client computing.

LINUX PRO: The thin client market for Linux has blossomed rapidly in the last two years. While Linux as an OS on the desktop makes up a small overall percentage of the market, it is gaining market share at a rate alarming to its competitors. This is perhaps all the more surprising because the world's number one supplier of thin client solutions, Wyse, didn't have a Linux client – until now. Recently, Wyse launched two thin client products for the European market, with a view to building up a share in this expanding market.

“In the last year, Linux thin client market grew 63 per cent in Europe – in three years time, Linux will be the biggest...”

FURTHER THIN CLIENT RESOURCES

WYSE – market leaders in what they like to call 'server centric' computing with thin clients. www.wyse.co.uk

LTSP – the linux Terminal Server Project is a system for using standard or stripped down desktops as diskless clients running X. www.ltsp.org

NEOWARE – One of the largest manufacturers of thin client devices, including the Linux based Capio device www.neoware.com

CITRIX SYSTEMS – Provide software and services for servers to run thin clients

from. The best known solution is the Metaframe range of software www.citrix.com

NIC – Linux based thin clients www.thenicstore.com/nic

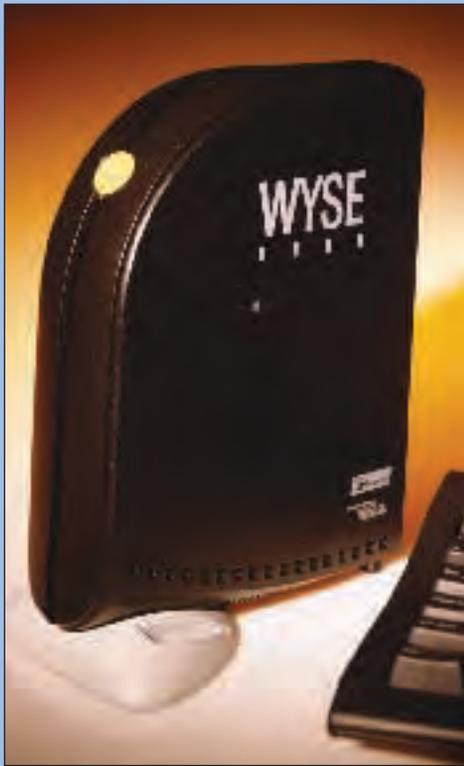
THINGUIN – more info on Etherboot and various other thin client/terminal projects www.thinguin.org

LIN:WARE – customised Linux client software for turning PCs into thin clients, supports the ICA protocol used by Citrix. Also supply thin clients. www.linware.com

processor, network connection, graphics output, memory and usually some sort of storage (preferably in the form of Flash RAM or similar) is desirable. It isn't beyond the realms of reason to 'repurpose' a standard desktop as a thin client. In fact, that's what projects such as LTSP were set up for (see boxout above). As processor speed is not often critical, this can be a useful retirement plan for desktops which have outlived their usefulness. Retro-fitting Flash storage can be a more complicated and costly business, so many projects work on the basis of using a completely diskless client that boots from the network itself. Unfortunately, this approach undermines some of the advantages we have already discussed – you won't necessarily be dealing with identical hardware, although in some ways administration becomes easier because the client machine is just an Xwindows terminal and doesn't include any software of its own.

WINTERM 5125SE – FEATURES

- Native X Windows, RDP, ICA protocols
- Netscape Navigator 4.77 browser with plugin support
- Built-in networking support for WiFi 802.11b wireless link
- Connects to any VESA-standard monitors
- Up-to-date Linux OS and components
- Remote management and configuration through Rapport™ for Wyse management software
- 10/100 BaseT Fast Ethernet, two serial, one parallel, three USB ports
- Configurable menu with application AutoStart and Restart
- Stateless, fanless, no batteries
- Read-only file system
- Three-year Buyer's Protection Plan



Linux is the place we have to be.

A lot of other companies we have contact with have similar stories in other markets and in other territories – Linux is where the growth is.

Also we have a lot of customers coming to us and specifically wanting an Open Source Linux thin client. We've not seen that before, but it's been happening a lot in the last 12 months.

LXP: Again that seems to be a common theme across the IT industry. Do you have any idea why

SY: A lot of it is cost. They like the lower cost of the Open Source software, and the freedom and openness it gives them. We may think of Linux in terms of the public sector, but the companies who have come to us are more the FTSE 100 type. I think the groundswell around Linux is happening everywhere.

“On Microsoft Windows CE, we can't offer Java – Linux changes that because the support for a JVM is much better.”

We see Linux as important to our growth and market share in Europe. The two products we have just announced are specifically for Europe. As a US-based company, I think that's a positive thing

LXP: Why have specific products for Europe?

SY: Europe is growing faster than the US in thin clients. A lot of that is because Europe is catching up. IDC predicted that by 2007, Western Europe will be as big as the US as a

thin client market, and our sales figure definitely support that. We do think that the biggest interest in Linux is coming from Europe. As I said this is an important market for us, so it makes sense to have a local product strategy. What do we see with Linux on the Desktop?

Linux has succeeded largely in the server market and has had little penetration on the desktop market. The reason is that people have tried to install Linux on PCs. In our opinion, that doesn't work so well because of issues of installation, support, drivers and user tampering. The thin client changes that. It's fundamentally different, and Linux is sitting on a solid-state device in Flash memory. It's installed in a factory or with a professional management tool. It basically makes the total cost of Linux drop down to way below that of Linux on a PC. In our opinion, it is the ideal delivery platform for Linux on a desktop. My own opinion is that Linux is going to arrive on the desktop in the shape of a thin client, not a PC

LXP: That's a very interesting point of view – there are obvious advantages. Even though the Linux distributions have improved the installation process by leaps and bounds, there are still issues there.

SY: The biggest issue I think. If you had 10,000 PCs and they were all the same, and you could create a byte-for-byte image of a Linux installation and copy it to all of them, it might be workable. 99 per cent of companies don't have that – they have a huge variety of old and new, different makes, different hardware needing different drivers – Linux isn't really designed for that environment.

LXP: Certainly it simplifies the management proposition...

SY: Greatly. And the other thing that thin clients can offer Linux is that, one of the things Linux has faced in the PC world is the lack of local applications, which is really a Win32 environment. The beauty of thin client is that you can access a Windows application running remotely on a server, or you can access a Java application, or a web-based application and so on, but you still have Linux on the desktop in a managed, low-cost way. It answers the criticism that Linux isn't able to access legacy applications.

Linux is very attractive for thin clients. Some people ask what Linux can do for thin clients – it has a low hardware requirement in terms of processor and memory, and it has excellent JVM support. On Microsoft Windows CE we can't offer Java; if customers want a JVM, they have traditionally bought a relatively more expensive XPE terminal. Linux changes that because the support for a JVM is much better. It also offers X Windows support, which is still an important market, especially in Europe.

The other thing is that you have all that useful Open Source software out there that you can embed into the Flash memory as required by different users.

LXP: Do many customers do that?

SY: It depends. The two products we have launched cater for both types of customer. The 5125 is the 'Ford Escort' of Linux terminals. It does 95 per cent of what the market asks of it, and it does it at a low price. It's £269+VAT. That



makes it the second-cheapest product in our whole range of more than 25 terminals. It offers Netscape, RDP, X Windows and terminal emulation. They are still important to some people.

That profile is really what the market ask for. If you want to change the software you can, but it has 32MB of Flash RAM and we really wouldn't recommend this terminal for customers wishing to experiment with Linux images. You take it out the box, plug it in and away you go. It ships with the *Rapport* management software. If you wanted to update version or remove software components and replace them with something else, you could do that with the management tool.

The 5455 is really more for the Linux power user. It's £379+VAT it has 64MB of Flash and 128MB of RAM. For a thin client, it has a fast processor at 550MHz, high-resolution video, it has a PCI slot. There is also a bay for internal drives if you really want that. The other difference is that this model offers a full Sun Java VM. This is more suitable for customers who want to run full java applications in the JVM. It would be suitable for customers who want to

Demonstration of a growing market: read *LXP41* to see how and why the University of Northumbria is replacing its desktop PCs with these Neoware thin client devices.

embed Java or Linux tools into the flash memory, which would typically be device drivers, or a small line of business application. For customers who want to add additional peripherals, an example would be a bank where they want dual monitor support for dealers.

LXP: Would you say that the target market for the Linux devices is the traditional markets that you have already been selling thin clients to in the past, or are these going to be likely to be more interesting in some sectors or others?

SY: Considering the price point we have pitched them at earlier in this interview, I think that interest will come from across the board. From a Linux point of view, that's interesting. In terms of the market, everyone talks about the public sector and education, but to be honest I think we're going to get interest from all sectors.

LXP: It seems that thin clients have come back into vogue. Do you see that continuing?

SY: Very much so. I worked at Oracle in 1994 when Ellison launched the 'Network computer'; Wyse actually launched the first thin client a little before that. In some ways he did the industry a bit of a disservice because it was a little oversold. It took the industry a number of years to get the price right, the support and complimentary software such as *Citrix* and so on to be stable and accepted. What we're seeing now is that in the last two years demand has really taken off. We are growing extremely rapidly, about 60 per cent year-on-year growth. We're picking up very large customers, household names – John Lewis, the Royal Mail, Virgin, a lot of NHS trusts – the whole ideology is being accepted by IT management. I would say 80 per cent of major organisations are piloting or investigating this type of technology.

If you look at the market penetration figures for thin clients in Europe, is about 2.5 per cent of the business desktop. In the UK it's probably higher – I would say about 3 or 4 per cent. If the growth rates continue; it's only going to take another few years before we get up to the 8-10 per cent level. When that happens, there will be a sort of tipping point where these devices become much more obvious and commonplace as businesses of many different types realise how they can gain from using Linux thin clients. ■

WINTERM 5455XL – FEATURES

- Native X Windows, RDP, ICA protocols
- Netscape Navigator 4.77 browser with JVM support
- Built-in networking support for WiFi 802.11b wireless
- Rugged metal box sits comfortably under monitor
- Updated Linux OS and components
- Remote management and configuration through *Rapport™* for Wyse management software
- 10/100 BaseT Fast Ethernet, one serial, one parallel, four USB, PCI slot
- Configurable menu with application AutoStart and Restart
- Flexible platform with many configuration choices
- Read-only file system
- Three-year Buyer's Protection Plan



STORAGE

SERVER

Sharq Bite: NAS redefined

Xinit's SharqStore devices make **PAUL HUDSON** excited about the future of Serial ATA storage...

INFORMATION

PRICE From £7985 + VAT
SUPPLIER Xinit Systems
PHONE +44 20 7538 8230
WEB www.xinitystems.com

SPECIFICATIONS

CPU 2x 3.06GHz Xeon
RAM 2GB ECC
HD 12x250GB SATA, RAID5
NETWORK Gigabit



Not so long ago, NAS was a long way from being ready for proper use in a corporate environment. Even now, many companies still rely on DLT tapes for backup and traditional servers for online storage. Now that Serial ATA – the new hard drive format designed to carry on accelerating IDE for the next decade or so – is generally available on new machines, it seems that the time of high-capacity, low-cost NAS devices might now have finally come.

The SharqStore

Heralded in our exclusive interview in *LXP47* with Iain Cheyne of Xinit Systems, the world of NAS has an all-new contender. Xinit's Sharq SPS234 scored 10/10 with a *Top Stuff* award in *LXF42*, with particularly high results in the hard drive category – this bodes very well for its NAS device.

The SharqStore devices come configured to your requirements, starting at 1U with four drives of 250GB each, making 750GB capacity in total, up to a 5U device with twenty-four drives of 250GB each, making five terabytes in total – our review machine had twelve drives, giving two terabytes in total. Note that each system comes pre-configured with RAID5 (although

you can of course request otherwise), so a fair slice of the space on each disk is given over to parity and these numbers take that into account – our machine, for example, is actually a 3GB machine once you discount parity. Each of the hard drives are connected up to a 3Ware 8500 SATA RAID controller, which should make each drive easily transfer up to 150MB/s.

To make sure the drives perform at their best, the machine is backed by twin 2.4GHz Xeon chips, 2GB RAM, and a gigabit network card – this would perform well enough as a standard server, but here it's likely to be eclipsed by the performance of the hard drives. The system has three power supply units to make sure the hard drives get all the power needed, and it will emit a warning beep if one of the power supplies is faulty or disconnected.

Openfiler storage

While the hardware inside is of top quality, the real advantage to this system lies in its software. Named *Openfiler*, Xinit has developed an all-new NAS management system that works entirely over the web using HTTPS, and which allows you to take control of all aspects of the machine through a point-and-click GUI. That's not to say the machine

can't be used as a standard server if you'd rather – it came pre-installed with a full version of Red Hat 8 and is perfectly usable. However, it was designed to be headless (sans monitor) and so the component of real importance, *Openfiler*, can be accessed from anywhere with network access and the right authentication information.

Using *Openfiler* you can partition the drives into volumes, create and delete directories on each volume, and configure each directory into shares accessible from your users. These users can be authenticated in a number of ways, although perhaps most popular will be the LDAP system – you can configure an LDAP server for the SharqStore to query, then use that to define which users can access shares on a fine-grained level.

Using the Services section of *Openfiler*, you can also configure how the system can be accessed – of most interest to Unix clients will be the NFS support (v2 and v3) as standard. Added to that, the drives can be accessed through SMB for people on Windows-dominant networks, and even WebDAV, which should technically be accessible from all computers that support the open standard. No matter which route is taken to the data, the same extensive authentication checks are still performed by the SharqStore, which makes the SharqStore family very much "set and forget" boxes.

Openfiler is itself built on top of a number of open source projects, and Xinit intends to open-source the whole system soon after SharqStore is released. By using Open Source

Three power supplies at the back mean that this machine gets all the juice it needs.



projects, and of course the choice of SATA as opposed to the more traditional SCSI devices, Xinit has managed to keep the overall cost of the system much lower than you would get elsewhere.

NA-speed

This box has twelve (yes, twelve) hard drives, of which RAID 5 dictates that a significant amount of each be dedicated to storing error-correction information. If any drives fail, RAID 5 will compensate automatically and the machine will keep on going – an absolute must for solid data reliability.

With the combination of the data being spread across so many drives, SATA, and a good-quality RAID controller, this machine should be able to handle any data storage requirement you want to throw at it. With such a powerful set of processors and 2GB of RAM also thrown in for good measure, this could easily have a place outside of the NAS market – anyone with a requirement for a solid-performing server and heaps of storage will find their needs easily met by this beast.

For use as a backup device, this naturally races ahead of other backup

competition such as tape drives – sure, it's not quite as portable as an eight-tape cartridge, but with gigabit networking as standard and a bargain price, you can afford to have two of these working in tandem with one mirroring the other.

Who needs 2TB?

At this price tag the machine certainly isn't a consumer device, so what kind of company needs this much space? Certainly any place that handles video processing will be well used to the hassle of constantly having to add more disk capacity, simply because disk space isn't increasing as fast as the demands of artists – one video render can easily chew up 500MB on its own, and when you have several working files in one project you're easily talking about 3–5GB per project. However, even that sizeable task is dwarfed by disc creation companies, where each final disc will take up 8GB of space alone, and each project is likely to top out at around 20GB.

Beyond the world of digital content creation, more and more companies are finding the sheer bulk of smaller files is slowly overwhelming them. News corporations such as Reuters and the

BBC go through hundreds of gigabytes of content every day, which make them particularly heavy users of NAS devices such as this. The BBC's recently announced *BBC Creative Archive*, a project where the Beeb hope to make available most if not all of their content online, has been estimated to require a massive 10 terabytes of disk space per channel per year, which makes this 2TB device seem like a tadpole! Of course, companies with archives as large as the BBC are few and far between, and, for the most part, two terabytes should be enough for pretty much everyone else.

This is a very capable machine even without the enormous hard drive capacity – Xinit has yet again managed to trump the competition with a device so uniquely easy-to-use that we reckon most people won't even realise it runs on Linux. The fact that it comes with a three-year, on-site, next-business-day warranty makes this is a safe buy for any company, and an inexpensive option no matter how you look at it. ■

VERDICT

XINIT SHARQSTORE

FOR

- + High capacity with quality SATA drives
- + Low price tag includes three-year on-site warranty

AGAINST

- Lack of SCSI drives keeps cost down at expense of speed
- 12-way SATA card used in a 16-bay box



Removal of hard drives for maintenance or replacement is easy – and RAID5 ensures that there's no letup in storage ability by switching to another drive if one becomes defunct.

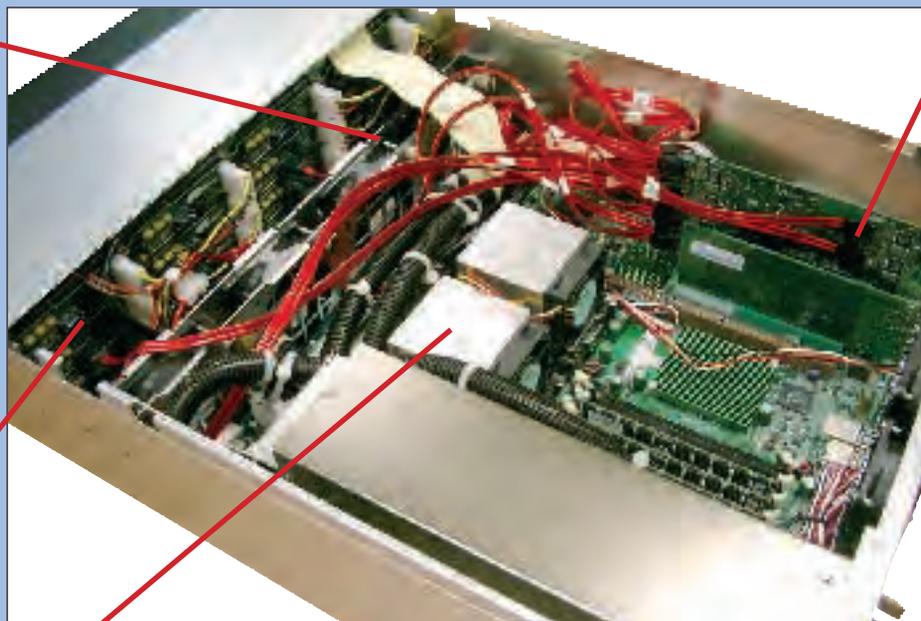
WHAT'S IN THE BOX?

FANS

Although hard drives don't generate that much heat compared to graphics cards, when you have twelve of them the heat accumulates quickly – there are quite a few fans in this machine, and it also draws a great deal of power to keep everything ticking over.

HARD DRIVES

This product is all about the hard drives, and with twelve of the latest SATA drives from Western Digital at its heart, it's not surprising that it performs like lightning.



SATA RAID CONTROLLER

In order to handle this many drives, and also because SATA hasn't really made it onboard just yet, the server comes packed with a 3Ware 12-way RAID controller. While this does guarantee lots of performance, it also means you'll need to change the RAID controller if you want to add more drives into the spare slots – a little extra hassle than necessary, perhaps.

CPUS

With two 2.4GHz CPUs behind this server, it can easily be used as a normal server – the Red Hat 8 installation was excellent, but Xinit will, as usual, install and configure either SUSE or Debian for you at the factory.

Rigorous testing with Nessus

Security scanning is one area that is often neglected in corporate security routines. This is a bit surprising, because if you get in a security consultant to harden servers or just test the system, running a scanner is about the first thing they will do, to get an overall picture of network security.

Perhaps the fear of performing regular scans is something to do with the assumption that scanning is something that unpleasant crackers do with malicious intent. But that's precisely why you should be doing it – you can bet that if you have a large enough web presence, someone else is already doing it to you!

Enter the *Nessus*

Nessus is a dual-purpose security analysis tool. As well as a simple portscanner, it will also perform simulated attacks, trying to exploit known vulnerabilities. This second stage is achieved using a simple plugin architecture. As new vulnerabilities become public knowledge, tests are written for them and distributed – this saves having to reinstall the software each time a new tests is written.

One of the real keys to the usefulness of *Nessus* is this plugin system. And by that we mean not so much the structure of the software, but the quality, availability and timeliness of the tests that are released.

NICK VEITCH
applies brute force
to the subject of
security testing.

Determine which tests you would like to run – obviously you may wish to omit inappropriate tests.

The tests themselves can cover almost any aspect of security that you can think of, from simple things like checking your SSH server doesn't respond to the vulnerable SSH1 protocol, or that you have changed the default passwords on various software, to more complicated checks such as installed sniffers and rootkits on already compromised servers.

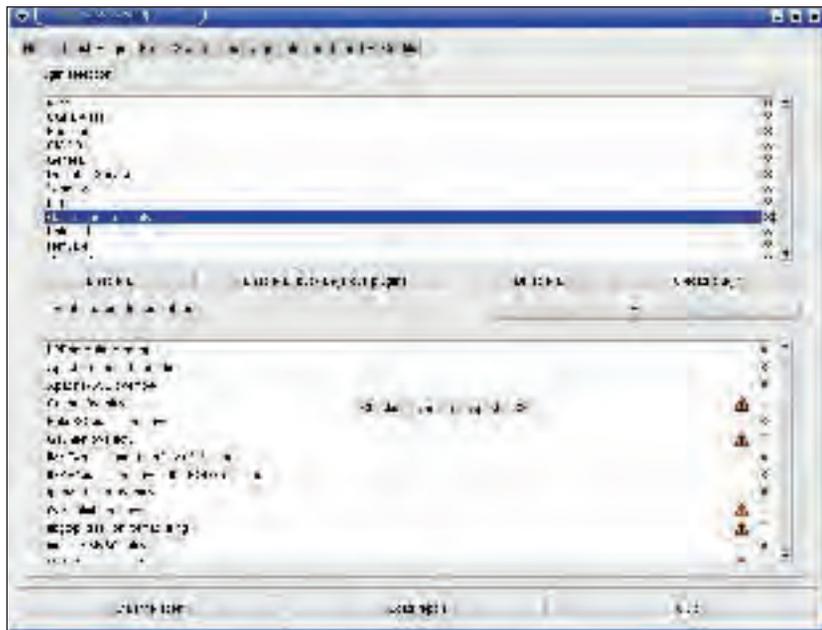
One important thing to realise is that some of these tests can force errors and 'undesired performance issues' in the target server if it fails the vulnerability test. You do not really want to run these on mission critical servers. Fortunately there is an option to disable all possible harmful tests.

As this is Open Source software, it is possible to create your own plugins for whatever flaws or systems you are concerned about. In practice, this is usually unnecessary, because the Open Source community is very good at providing these for almost anything you can think of. We should point out here that this isn't restricted to Linux/BSD servers, but covers, potentially more extensively, Windows-based servers too, including specific services and vulnerabilities that are only relevant to that platform.

Using *Nessus*

Once you have obtained and installed the *Nessus* software (see the box opposite), there are only a few more steps before you can start probing hosts.

You can create a configuration file for the server, but this is actually most easily done by running the client – it will



KEEP IT UP TO DATE!

OBVIOUSLY, FOR A SECURITY SCANNER TO REMAIN relevant, you must keep up with the vulnerabilities and security holes that become apparent. Plugins are updated and new ones released regularly, and are listed on the main *Nessus* site, but you don't want to resort to installing updates manually each day.

The *Nessus* distribution includes an update script, which will fetch and install any updates from the main *Nessus* site. It only makes use of lynx, tar and gzip, which come with most Linux distros, so you are bound to have them on your system anyway.

Obviously, the preferred method is to run the script via a cron job on a regular basis to make sure everything is up-to-date.

create a default configuration for you. What you must do before running the client is create a user account. *Nessus* maintains its own list of users separately from the system. A utility, *nessus-adduser* will allow you to add users to the system and verify with passwords or certificates depending on which is more appropriate for you.

For added security, if more than one person is going to be using *Nessus*, it is recommended to edit the configuration file and add restrictions to the user accounts. These restrictions can include restrictions to limit the servers that can be tested by any particular user. The documentation and the man page for the *adduser* tool explain the syntax further.

Running the client brings up a GUI where you can select the tests to be performed and the targets you wish to scan. The Plugins tab will enable you to select the specific plugins you wish to run. The plugins are grouped together into categories to make it easier to perform specific tests. A good place to start is by selecting 'Enable all but dangerous plugins'; then deselect any other tests you don't want (unless of course you wish to try the potentially harmful tests).

There are options to peruse, but really you are best off reading the supplied documentation to determine which ones best suit your purposes. When you are happy with the setup, proceed to the 'Target' tab to select the targets for your tests. You can supply a single host or a comma-separated list. Click on the 'Start the scan' button at the bottom to begin. The tests may take some time!

Understanding the output

The report screen includes five panes of information. These are user configurable, to make it easier to list messages the way you'd prefer – by type, by server or by the relevant port number, for instance. Clicking on the relevant entries will bring up information in the large report pane, specifying exactly why this item has been flagged.

There are three levels of severity – a note, a warning, and a detected security hole. Notes contain information that may be relevant to any problem that has been encountered. Pretty much any scan will produce some notes, as these will include simple information such as "An SMTP server is running on this port", or possibly more important information such as "An unknown service is running on this port, which is usually reserved for ftp".

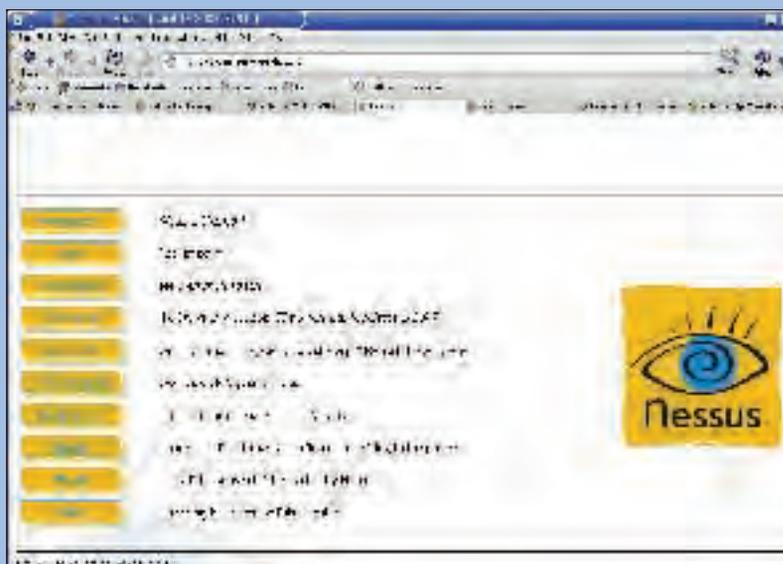
Warnings contain low risk factor comments on the server setup, including things like SMTP servers allowing relaying (this may not necessarily be an exploitable state, but is worth making sure about), or just useful comments. E.g., if your webserver includes a robots.txt file you will get a note warning you of the dangers of including sensitive directories in the restricted list.

Security hole messages are reserved for serious problems, such as out-of-date packages and recognised exploits. If you encounter any of these, they should really be attended to as a matter of urgency.

Helpfully, the messages not only outline the security flaw, but more often than not give some helpful advice about how to fix the problem. Short of downloading the patches and installing them itself, it's difficult to imagine how the client could be any more helpful! ■

GETTING AND INSTALLING NESSUS

Make sure you are using the latest version



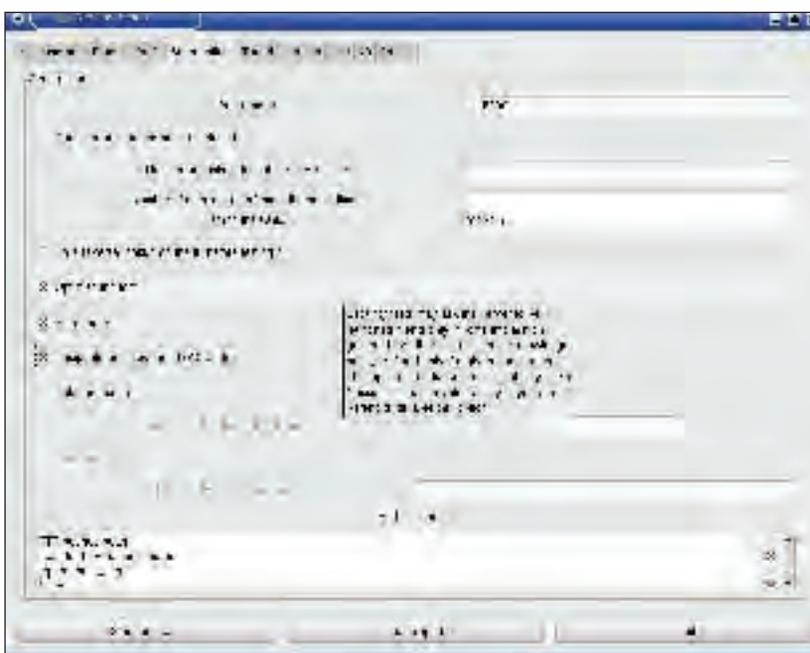
AS WITH MOST SECURITY TOOLS, YOU are best off going to the main website and making sure you get the very latest version available – having an out-of-date security tool is worse than using none at all, as you may be fooled into thinking your servers are fine!

The main site at www.nessus.org houses the source for the software. Possibly the easiest way to do this is using the autoinstaller, which is available at the site. Download and run this script to automatically fetch the latest packages and install them. Source code tarballs are also available here.

Because of the popularity of *Nessus*, you will probably find a distro-specific version in the usual places. There are certainly debian, SUSE, Red Hat and Mandrake packages available.

Nessus works as a client/server pair. The server does the actual work, the client provides a GUI for specifying the parameters and displaying the output. This approach works well, as you can more easily position the 'source' of the attack in terms of your network.

Clients are available for Windows too, so you don't need a Linux/Unix desktop to run the software if you don't want it.



Tooltips in the GUI client can give you some advice on the available options.

HOSTING

WEB HOST FOCUS

Rackspace Managed Hosting

Famous for its creed of “fanatical support”, Rackspace came first in our web host Roundup in September, scooping our *Editor’s Choice* award. One of their biggest advantages, though, is also quite a well-kept secret for some reason – Rackspace offers a 99.999% uptime service level agreement, which means it guarantee your server will experience less than five minutes of downtime *every year*. It’s a bold claim, but clearly one the company sticks to, as it has over 7,000 customers in more than 80 countries!

Not content with providing the highest SLA we’ve seen, Rackspace has recently announced its intention to provide a 100% uptime guarantee – that is, your server will *never* be down due to bandwidth issues, and, if for some reason it *does* go down, customers can claim back up to 100% of the monthly server fee. Combined with the extra uptime

In the second of our series taking a detailed look at individual web hosts, PAUL HUDSON checks out the company that won our comparison in September’s issue.

guarantee, Rackspace also recently announced a one-hour turnaround time in the case of hardware failure, which means that it promises to replace any hardware in your server that fails within one hour or, again, you can claim back up to 100% of the monthly server fee – a bullet-proof combination of guarantees that should make any sysadmin wonder why they should host their site themselves.

In order to provide the high quality of support that people expect, Rackspace has worked hard to make sure there are teams of experts available for all customers so that everyone with a Rackspace account has access to 24x7 support for all their hardware and core software issues. In our tests we found that Rackspace answered its support telephone number immediately, with no answering machines, hold music, or automated computer systems.

DOMINIC MONKHOUSE SPEAKS

The MD of Rackspace Managed Hosting Europe talks to *Linux Pro* about the company, how Rackspace uses Linux, and plans for the future

LINUX PRO: WHAT ARE RACKSPACE’S UNIQUE advantages over its competitors?

DOMINIC MONKHOUSE: I’d say that it’s our people that sets us apart from other hosting companies. Offering 24/7 customer service is key, but at the end of the day it is the quality and commitment of the people who manage and support a site that counts. Rackspace has developed Fanatical Support™, which is a 100% people-focused service. In our mission to deliver this we have absolutely no automated call handling – a person, not a machine, answers every technical support call. Try phoning us at 3:00am to find out for yourself!

Recent research shows that 67 per cent of customers who come to Rackspace from other hosting providers do so because their previous hosting company let them down. It’s no longer a case of “better the devil you know”, because people are wising up to the fact that good



“We use both Intel- and AMD-powered servers... we employ what fits best with our customers’ existing IT strategies to make sure that they get what they expect”

DOMINIC MONKHOUSE, RACKSPACE

customer service is essential – and they’re coming to Rackspace to find it.

LXP: What is Rackspace’s key target market?

DM: Any organisation that finds the Internet critical

to its business is a potential prospect. Our customers range from small independent web developers and system integrators to Fortune 1000 companies who need complex solutions. What they all have in common is that they are serious about

For range of products, Rackspace has something for just about everyone. Their lowest-cost system starts at £99 a month, for which you can get a 1.5GHz CPU with 256MB RAM, 40GB of hard disk space, and 30GB of burstable bandwidth, all pre-installed with Red Hat Linux ES 2.1. At the other end of the scale, services that cost upwards of £20,000 a month can include load-balancers, firewalls, clustering, and other enterprise-level technology. It's worth noting that Rackspace use Red Hat's Enterprise Linux distros, which, while probably quite costly for Rackspace, ensures 100% system compatibility and the highest levels of reliability.

With large data centres in both San Antonio and London, Rackspace aim to provide redundancy across the board – power, connectivity, and HVAC are all backed up so well that Rackspace has had 100% uptime for over 30 months now, a feat that most other ISPs would balk at.

Testing Rackspace

In our September feature, Rackspace came first by providing far and away the best product in our price bracket, and backed that up with high-quality support and first-class performance. We were also very impressed by the MyRackspace online functionality, which contains help, tools, and online support information for each account, even down to keeping a log of all support calls to Rackspace and its solutions. Of course, perhaps the acid test of what we think of Rackspace is the fact that when it came to finding a guru to answer our Linux system administration questions, we looked no further than Hans Huberland, one of the many Linux experts working in Rackspace Europe's tech support!

Rackspace and The Register

When it comes to online IT publications, *The Register* pretty much leads the way with over 2.1 million readers around the world, accessing a total over over 23 million pages every month; as such, it has phenomenally high technology



In old-fashioned connectivity terms, West Drayton is handy for Heathrow, London and the M25, but this hi-tech data centre hosts sites for many businesses across the world.

requirements. As *The Register's* first ever technology partner, Rackspace was brought in to take control of all *The Register's* hosting requirements, providing *The Register* with a total of four servers running Linux, two of which are web servers and two are database servers, but this solution was designed with scalability in mind – more servers can be easily added as the site's requirements grow. All four servers are at Rackspace's fully secure, state-of-the-art Data Centre in West Drayton, near Heathrow.

Linus Birtles, managing director at *The Register*, said: "This is a first for *The Register*. It's been a long and difficult journey to find the company that would suit our unique needs. But with Rackspace we found a cultural and technical fit. We are confident that Rackspace will enable us to serve our growing readership in the manner that it is accustomed"

A spokesperson at Rackspace added: "Nearly all of us in the Rackspace office use *The Register* as our home page, so we are delighted to have it on board! *The Register* was very aware of its hosting needs and it was the combination of our people, our experience with Linux and our knowledge of hosting a complete business infrastructure that secured us this prestigious customer." ■

their web sites and need reliable hosting solutions. Feedback from our customers shows that 95 per cent of them would recommend us to others, which is a figure that really speaks for itself.

LXP: What kind of hardware does Rackspace deploy, and what are your hardware plans for the future?

DM: Because Rackspace builds hosting solutions to meet individual customers' needs, we use a range of hardware including Intel- and AMD-powered servers. We employ what fits best with our customers' existing IT strategies to make sure they get what they expect of us. We're actually the largest user of AMD-powered servers in Europe.

Rackspace's network has been designated by Cisco Systems as "Cisco-powered", meaning that it has passed Cisco's compliance audit for security, redundancy, and speed. In fact, Rackspace is one of only a handful of hosting companies in the world to

earn the designation "Cisco-powered"

LXP: How much does Linux figure into your hosting plans?

DM: When Rackspace was founded four years ago, we forged a partnership with Red Hat and we're a staunch supporter of the Open Source community. Rackspace's internal systems are all run on Linux, and its technicians are qualified Linux experts. Half of our new business is Linux-based, as is 60% of our total business.

Rackspace offers the latest Linux technology – Red Hat Enterprise Linux AS and Red Hat Enterprise Linux ES, which are designed for mission-critical enterprise computing and certified by top enterprise software vendors to cover the full range of commercial IT environments, from edge-of-network to enterprise. I'd say that our long-standing partnership with Red Hat is what gives us that extra edge.

LXP: Where do you see the hosting market being in three years time?

DM: The hosting market is still very fragmented – I'd estimate that the top 10 providers only have about 15% market-share. Consolidation is therefore inevitable in this environment. The make or break factor for hosting providers is their ability to stay ahead of the field and understand what customers really want. Rackspace is currently experiencing 10% growth month on month and our success owes much to our quality of service and our satisfied customers. We see some clear winners emerging in the market over the next three years, of which we are confident that Rackspace will be one.

In terms of Linux, we see many opportunities ahead for the hosting industry. The open source community is growing and organisations are starting to re-deploy Linux as they move over from Sun and Solaris platforms. Companies are moving because of the stability and processing power Linux offers.