

LXF49
Linux Format January 2004
● JAVA DESKTOP ● HP WORKSTATIONS ● RED HAT ENTERPRISE SERVER 3 ● USING CSS ● AFPS ● DIGITAL RIGHTS ● DVD PLAYERS ● MAYA 5 ● EPIA CL10000

■ HOTTEST NEWS ■ LATEST REVIEWS ■ BEST TUTORIALS ■ YOUR PROBLEMS SOLVED

LINUX

THE UK'S BEST-SELLING LINUX MAGAZINE!

FORMAT



JAVA

ON THE DESKTOP

Sun's desktop strategy revealed – **EXCLUSIVE** details of new 'Java Desktop' distribution

JAVA

STANDING UP FOR DIGITAL RIGHTS!

The great 'broken CD' scandal and how it affects you **p52**



WHO ARE THE AFFS?

Find out who campaigns for Free Software and how you can help **p56**

HP WORKSTATIONS



Desktops that pack a punch and won't break the budget - the latest Linux hardware from HP on test inside **p18**

TOP PLAYERS

The latest DVD player software rounded up and rated **p34**

RED HAT ENTERPRISE

New software and a new strategy for RH Enterprise Server – reviewed in-depth **p22**

GET MORE FROM LINUX

- Master the command line **p62**
- Using CSS on your website **p66**
- Using Blender meshes **p70**
- Finding the best GIMP filter **p74**
- Get to grips with PHP5 **p78**
- Understand networking **p86**

'Blue Gene will revolutionise the way supercomputers and servers are built', William Pulleyblank, IBM **p07**

DVD issue also available Printed in the UK
LXF49 JANUARY 2004 £5.99

 Media with passion 9 771470 423019 0 1
 (£5.99 outside UK and Republic of Ireland)

www.linuxformat.co.uk

The Sunny side of Linux

Sun has certainly had an unusual relationship with Linux over the years. First the company was all over it, then to effect some notion of indifference, it derided Linux as useful, but not for 'serious' computing. Now it is passionate again, and this time with feeling. Linux forms a major part of the new Java Desktop strategy – a confluence of all the goals Sun is now focusing on. The name actually suggests that it is an excellent platform on which to run Java applications, rather than it being built with Java. In fact, under the bonnet and behind the gloriously redesigned splash screens, lurks a slightly modified version of SUSE Linux. If the network is the computer, and Sun is the network, then it seems Linux is the OS. Anyway, Paul Hudson went to Berlin's launch event to decipher the meat from the message – enjoy his unravellings starting on page 46.

Recent changes in copyright law (you mean you didn't notice? No, neither did anyone else) have brought the issues of intellectual property

rights to the fore. It would be hypocritical to suggest that as a magazine we are not in favour of legal management of intellectual property rights, but sections of the European Union Copyright Directive seem like just another tool to erode the freedoms of the consumer and line the pockets of protectionist media entities. How can purposefully ignoring CD standards and creating discs with flaws added to them on purpose be for the common good? Is deliberately making an audio CD unreadable on your Linux box really an effective weapon against copyright theft? You'll find out more about audio CD protection in our second feature, kicking off on page 52.

In a coincidental manoeuvre (honest), our *What On Earth* feature this issue explores the workings and purpose of the AFFS, an organisation whose *raison d'être* is the safeguarding of your rights and the promotion of Open standards and Free Open Source Software. Find out what they get up to and how you can get involved on page 56. Don't forget!




Nick Veitch EDITOR

Sun – serious about Linux, Java and the desktop p46

Would you pay £12.99 for a broken CD? You probably already have p52

Standing up for FOSS in the UK – discover the AFFS p56



MEET LINUX FORMAT'S TEAM OF WRITERS...



Andrew Channelle
The Linux beginners' best friend, Andy is also our roving reporter sniffing out stories for the News section.

Tom Wilkinson
Knows *Apache*, *httpd*, *BIND*, *Sendmail*, *Exim*, *PHP*, *FreeVSD*, *MySQL*, and *TCP/IP*... but only during darkness hours.



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

Hoyt Duff
Fishing pier proprietor and 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.

Maurice Kelly
Electronic engineer, coder and music fan who *knows* that history will prove the EUCD totally unworkable.



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

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



Paul Hudson
Went to Berlin right on deadline day for the launch of the Java Desktop, and all we got was this lousy T-shirt...

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



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

Newsdesk

- SGI's Linux supercomputer ● Revolution-OS Linux movie is out on DVD ● Debian Woody ● DLC
- Linux in Brazil ● Red Hat treason? ● Sun million desktops ● Motorola Linux PDA/phone ● Wine Rack

LINUX MARKET CONSOLIDATION

Novell's SUSE purchase gets positive reaction



Novell's CEO Jack Messman announces more good news for Linux.

With the ink not yet dry on Novell's statement of intent to purchase SUSE Linux for US\$210 million, speculation has begun on the next big takeover. And Novell itself may well be in the firing line for a takeover by an even bigger fish...

According to a recent story on a Columbian news website, insiders at IBM have suggested that Big Blue may make a bid for Novell within a month:

that's by the time you read this issue of *Linux Format!*.

IBM has enjoyed a very tight relationship with SUSE, with the companies jointly bidding and winning the famous Munich desktop deal, which will see 14,000 Munich city workers sitting down at Linux desktops over the next few years. The company also invested US\$50 million in Novell at the news of the SUSE acquisition. Of course, this is just speculation...

Gateway Computers, which dabbled in Linux in the server space a few years ago before scaling back, has announced that it will soon begin shipping servers built around SUSE Linux Enterprise Server. The packages will include support for patch and driver updates and one year's maintenance. The company said the alliance would offer a single point of purchase, support and maintenance for SUSE Linux Standard Server 8 and

SUSE LINUX Enterprise Server 8. Scott Weinbrandt, senior vice-president and general manager of Gateway's Enterprise Systems Division, said that supporting Linux was a logical step for the company, as businesses of all sizes were increasingly looking to the security, flexibility and cost advantages of Linux-based solutions to address IT needs throughout their organisations. "SUSE's status as the first Linux

LINUX SUPERCOMPUTER

IBM's true colours



Will Blue Gene be an eponym of 21st century computing?

Blue Gene, IBM's new US\$100 million supercomputer project, will forsake the company's own version of UNIX in favour of Linux – a move which represents an act of faith in Open Source software and the development model that drives it. The goal of Blue Gene is to build a new breed of supercomputer capable of crunching one quadrillion (one petaflop) calculations per second. The first systems should be available in late 2004/early 2005.

IBM says the first production model, Blue Gene/L, will house 65,000 processors, 150,000,000MB of memory, and will be used by the Lawrence Livermore National Laboratory in the US to simulate nuclear explosions. A test version of the system – which only occupies about 1 cubic metre of rack space – recently landed at number 73 in the list of the world's most powerful computers. The full model is estimated to be 128 times

larger and six times faster than today's fastest supercomputers.

"Blue Gene's entry onto the Top500 list marks a fracture in the history of supercomputing," said William Pulleyblank, director of exploratory server systems, IBM Research. "It will revolutionise the way supercomputers and servers are built and broaden the kinds of applications we can run on them."

Pulleyblank said the decision to adopt Linux for the supercomputer was made partly thanks to the size and dedication of the Open Source community, and that building a new OS within IBM would take an unfeasible amount of time and money. "We chose Linux because it's open and because we believed it could be extended to run a computer the size of Blue Gene." He said there was considerable advantage in using an operating system supported by the Open Source community.

3D MODELLING

Blender by the book

To complement the latest *Blender* 3D release, the project has produced a 600-page user manual featuring tips, examples and a full-colour 16-page image gallery. The organisation has begun a pre-sales campaign in an attempt to bridge a financial gap.

Sales of the manual have previously been instrumental in the process of keeping *Blender* alive, project leaders say, and the new one is no different. In a statement, Ton Roosendaal said a sales target of 500 had been set to cover expenses and printing and that an additional 500 sales would "allow us to close the books on a positive result." The book, scheduled for release in the last



Make the most of Blender...

week of December, is available for 44.50 Euros, though pre-orders qualify for a discounted price of 35.00 Euros, and the first 250 will be signed by the author. You could even win your own copy in LXF50! www.blender3d.org/

distributor to achieve the federal Common Criteria standard for security is clear proof of the company's strength as an enterprise market partner."

Holger Dyroff, SUSE's American head, said this was another "another significant milestone" in the snowballing adoption of Linux as an enterprise-ready operating system solution around the world.

"Gateway systems and support, coupled with SUSE software, provides customers with a true win-win situation. And of course Gateway will continue to serve as a single point of accountability for customers that migrate or choose the SUSE platform on Gateway."

Not happy

SCO predictably marked news of Novell's purchase with a fresh press offensive, in which Darl McBride claimed that the deal, if it goes ahead, would break a non-compete agreement reached between the two companies back in 1995, when SCO bought UNIX System V.

"We bought the Unix System 5 rights from Novell back in 1995, and there was non-compete language that would prevent Novell from competing against our core offerings," McBride blustered during a press conference. "Linux is a knock-off of Unix. There can't be a more straightforward reading of the non-compete clause." He stated that SCO would take steps to enforce the non-compete clause if and when the deal went ahead.

A responding statement on the issue from Novell said "Mr McBride's characterisation of the agreements between Novell and SCO is inaccurate. There is no non-compete provision in those contracts, and the pending acquisition of SUSE Linux does not violate any agreement between Novell and SCO." The statement added that no official notice had been received.

Jono Bacon

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



COMMENT

Simple is as simple does

“ Since the penguin-flavoured Linux first seated itself on my hard disk back in the dark ages, I've been privileged to be privy to its development and maturity. As Linux has outgrown its metaphorical goatee (after a period of staying out until late and drinking too much) it is turning into a mature system that can be used for all kinds of purposes. Flexibility is the key, and Linux (who looks so much like a pre-pubescent Eminem in the IBM ad) has grown up.

Though we've the burgeoning IBM, HP, Novell, SGI, Oracle and countless other suit-laden corporates behind us, Linux has still managed to remain a system that is essentially for everyone. There is a distinct humility knowing that the OS on my computer is ours, and no one can take it from us. Yes, that includes you, SCO.

There is still indeed work to be done. You know that somewhere, someone wants to run Linux on a computer with a gazillion processors and a zibabyte of RAM and this kind of support will continue. Linux will continue to run on PCs, PDAs, toasters, Christmas trees and all manner of other crazy platforms. Hardly a day goes past without some bright spark managing to contort the OS to run on something even more unusual.

The key to this whole scenario is simplicity. The OS is free, the code is available and the bright minds are out there. It just goes to show what people can do when they are given the opportunity to contribute to something that is not just theirs, but everyone's.

We live in exciting times for Linux – and computing as a whole – and who knows what's next in store for 2004...? ”

LINUX SUPERCOMPUTER II

SGL pushing the envelope

SGI used the SC2003 Conference in Phoenix, Arizona, to demonstrate a new Altix 3000 supercomputer powered by 128 Intel Itanium 2 processors and the company's own NUMAflex shared memory technology. Jim Taft, of

NASA's Ames Research facility, presented a paper on the creation of the first 512-processor Altix-based single system, and Jon 'maddog' Hall detailed the changes taking place within the Linux kernel to make it the best choice High Performance Computing (HPC) applications such



The Ames Research Centre is at the cutting edge of HPC.



SGL claims NASA's new machine is the largest Linux supercomputer.

as clusters, NUMA machines and emerging technologies like The Grid.

Bob Bishop, chairman and CEO of SGI, said the company had made great strides since the industry conference last year. "SGI has introduced significant advancements across all product lines as a continuation of its total focus on the high-performance technical marketplace," he said.

SGI also unveiled the 'World's Largest Supercomputer' which has been installed at the Ames research labs to monitor and model changes in temperatures in the World's oceans. Within weeks of "attaining record levels of sustained performance and scalability on 256-processors," the Ames team doubled the size of the system to create "by far the largest supercomputer ever to run on the Linux OS."

Linux Web Watch/



Download huge 12MB pictures.



Download pictures from history.



Find out about your fave movies.



The history of Linux on your TV!

Flash! Bang! Wallop! What a picture!

With the widespread adoption of broadband continuing apace, browsing huge picture resources is even easier...

To mark the launch of the Cinepaint Digital Film Library (www.ibiblio.org/cinepaint/), LXF spent quite a few happy hours perusing other cinematic resources on the Internet. The Cinepaint library – which takes its name from the Linux film paint application formerly known as *FilmGIMP* – is the "world's first online archive of cinema grade moving imagery." Each frame of

each film weighs in at a hefty 12MB. The selection is extremely limited at present, but it's a great idea that is worth keeping an eye on.

The Pathe Archives, available at www.pathearchives.com/welcome-en.asp, features an enormous collection of archive footage and photographs dating back to 1896. Though registration is required, this is

an essential site for anyone interested in history as it has been presented through the media.

For more up-to-date info, the Internet Movie Database (www.uk.imdb.com) is a long-standing favourite. As well as a précis and cast list of almost every film you'd care to mention, there are links to reviews of current fare, fan sites and DVD, video and merchandise sellers.

Finally, Linux – The Movie, aka Revolution-OS, has a dedicated website (www.revolution-os.com) which brings together photos, links, information about screens and the all-important opportunity to buy the DVD of the independent film which charts the rise of our favourite independent OS. Everyone you'd expect to see is in it – plus more besides!

NEWSBYTES

■ **Microsoft** has announced the latest version of Virtual PC software, the product it purchased from Connectix. Despite fears to the contrary, the virtualisation software remains Linux-compatible, though Microsoft says it is not an 'officially supported' guest OS.



■ After the Mac equals good-guy, Windows equals bad-guy theories which have surrounded the TV series **24**, the producers have given Jack Bauer a new weapon in the fight against global terrorism: a PC running Linux and KDE. Unfortunately for Jack he appears to be running KDE 1. How can he manage without alpha-blended icons and drop-shadowed windows?

■ **Kaspersky Labs** has updated its virus scanning software for Linux- and UNIX-based mail servers. Version 5, which can be managed via *Webmin*, now includes the ability to monitor email data streams in real-time and, the developer says, features an 'anti-virus kernel' which takes the pressure off the server's CPU. Purchaser of version 4 can download the update free of charge from www.kaspersky.com.

■ **Netcraft's** latest survey suggests that *Apache's* lead in the web server stakes is growing. Monthly figures show *Apache's* popularity has jumped by 2.8 per cent and now accounts for over two-thirds of the Internet.

■ **Red Hat's** Matthew Szulik has said that, while Linux is more than capable in an enterprise environment, home users were better off sticking with Windows for the foreseeable future. Szulik said that potential users might hit a stumbling block (most commonly with device drivers) and be put off in the future by this negative experience. "It is my view that the technology needs to mature a bit more," he said.

■ The **government of Brazil** has become the latest to adopt a nationwide Linux/IT strategy. The government has begun rolling out a public Internet access campaign in Sao Paulo using Linux-based terminals, and will soon spread the initiative nationwide. Meanwhile, the Spanish government has launched a program to put Linux on the desktops of 80,000 students. The deal – struck with Novell – is expected to yield annual savings of 20 million Euros per year for the country's education department.



This spoof of Debian Linux could end up giving some Linux users a completely different idea of the kind of Woody we're talking about...

DEBIAN DEVELOPMENT

Debian Woody reaches for the skies...

The 'stable' version of Debian

Woody is nearing with the launch of a second release candidate "which mainly adds security updates ... along with a few corrections of serious bugs." The release note suggested that those who regularly update their systems from security.debian.org will have most of the patches already.

Five packages, including a spell and *Rocks 'n' Diamonds* had to be removed from the distribution recently due to licensing issues.

The release was delayed for a few days due to an attempted attack on a number of Debian servers. An announcement said that *master* (Bug Tracking System), *murphy* (mailing lists), *gluck* (web, cvs, people) and *klecker* (security, non-us, web search, www-master, qa) were affected by the incident, but that the archive itself was not touched. However, extensive measures were taken to verify the state of the security archive before it was made available again.

WHERE IS THE LOVE...?

Standing room only on the Linux bandwagon

Ransom Love, former CEO of SCO (nee Caldera), has returned to the Linux fold after a 12-month hiatus, joining the board of directors of Progeny, the bespoke Debian reseller. In an interview with Cnet, Love – who left Novell to set up Caldera in 1994 – said it felt like the right time to "jump back" into the market, especially with the Novell purchase of SUSE. "When we introduced Linux seven or eight years ago, we knew Novell needed Linux. I've actually been wondering why it's taken them so long, but certain things needed to be there." He added

that Novell would bring "stability and strength" to the application infrastructure, as well as providing a "buffer" between SCO and the industry. "I'm almost certain that Novell has existing rights for using Unix products, so they may very well be indemnified. When they sold Unix to SCO, they kept a lot of stuff themselves." The Novell purchase, he said, may define the industry for the coming years and that Progeny – with its strong links to the Open Source community – was in a "strong position" to exploit the rise of Linux in the Enterprise sector.

Hoyt Duff

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



COMMENT

WTF wrestling

“ In a aberrant display of political correctness, a municipal administrator in California objected to the terms 'master/slave' as applied to hard drives and asked vendors to replace all such offensive terms on all equipment. What?”

This is reminiscent of occasional debates on Linux mailing lists about cursing in the source code (the kernel is famous for this) and terms like 'fork child', 'kill parent', and 'spawn daemon'. As we know, many of these terms are jargon to which only the uninformed and unenlightened take umbrage.

While many consider free speech to be an absolute right, it's a right that's abused when individuals feel pervasively compelled to say anything, anywhere, anytime. As a practical matter, some speech is restricted as a matter of law for the public good, and as a matter of custom for the mores of society, aka 'good taste'. While individuals who use the cover of free speech to deliberately violate good taste should be admonished and shunned by those they inflict their behaviour upon, it's also incumbent upon those so offended to take steps to avoid these annoying people rather than forcing them to behave in some 'acceptable' manner.

While cursing in the code may be humorous (and even therapeutic), some self-restraint is in order to use it only when necessary, as it is, by definition, outrageous. If you are offended by it, by all means admonish the author (expecting to be admonished in return and not expecting curses to be removed), then choose to ignore the cursing or stop using the code if it matters that much. We just shake our heads in disbelief... ”

LINUX IN CHINA

Sun: million desktop deal

Sun Microsystems is touting a deal made with a consortium of Chinese tech companies which could see the Linux-based Java Desktop System installed on 200 million desktops in People's Republic of China. The consortium, the China Standard Software Company (CSSC) has the task of 'bridging the digital divide' for 1.3 billion citizens.

The deal will bring a CSSC-branded edition of Sun's Java Desktop System to up to one million people per year.

Han Naiping, CSSC General Manager, said the initiative would advance the organisation's technology strategy and bring a fully localised operating system to millions of people. "With the Java Desktop System, Sun will provide the necessary technology to significantly strengthen our desktop initiatives," He said. "We expect to continue evolving our collaboration to cover a wider range of partner projects in the future."

The Chinese Government has skirted around officially supporting Open Source for years, but recently begun to 'firm up' its stance. Government spokesman Li Wuqiang said the benefits of Open Source were becoming apparent, not just in China, but all around the world. "A desktop solution based on Open



CEO, President and Chairman Scott McNealy prepares Sun for a huge portion of Chinese takeaway.

Cross-platform computing is likely to be the lynchpin of future development.

standards means more choice, an affordable price and a higher level of information security," he said.

In September, the governments of Japan, South Korea and China unveiled plans to develop an Open Source

alternative to Windows, with Linux as the widely tipped starting point. A Chinese Ministry of Information spokesman said the government would be funding Linux development as a way of enhancing growth in China's burgeoning IT industry.

Embedded Linux News

- **MontaVista** has 'Open Sourced' a number of its development technologies, and established community access sites on Sourceforge. The latest contributions include Dynamic Power Management, to reduce power consumption; High Resolution POSIX Timers to enhance real-time performance; Variable Scheduling Timeouts (VST), to extend the stand-by time of battery powered devices; and Protected RAM File System (PRAMFS) to protect data. In the near future, MontaVista plans to add Prioritised Work Queues (PWQ) technology to the list.

Spokesman Kevin Morgan said the releases of these applications were a call to arms. "MontaVista Software invites our open standards and Open Source community peers to join forces with us and with industry-leading hardware vendors to implement standard technology on the broadest possible platform base," he enthused.

- **Elematics**, developers of Multi-carrier network control plane software is bringing its Intelligent Network Control Plane (INCP)

software to MontaVista's Linux Carrier Grade Edition (CGE) on Intel processor. The company says the move towards Linux is in response to customer demand.

- **Motorola's** latest Linux+Java phone, the A760 has been announced. Combining phone, PDA, camera and MP3 facilities, the A760 is aimed at corporate users who need constant access to email when out of the office. The phone also includes a full Bluetooth implementation and 65,000 colour screen.

At last – a phone/PDA hybrid that actually works?



SCO News

- Executives claim that a bill will be sent out to a 'prominent Linux user' very soon. Meanwhile the deadline for purchasing a Linux IP licence at the bargain price of \$699 per CPU has been extended again. One reading of SCO's latest SEC filings suggest the licensing program is not going as well as expected. The filing reveals that Microsoft has made at least two large payments to SCO, the last one in the sum of US\$8 million.

- In response to SCO's issuing of subpoenas to prominent Linux developers including Linus Torvalds and Richard Stallman, IBM has called for access to documents pertaining to the recent \$50 million investment by BayStar and the 'buy' advice issued by Deutsch Bank.

- The various BSDs will be in SCO's line-of-sight if the Linux court battle goes its way, CEO Darl McBride has said.

- The litigation doesn't seem to be adversely affecting growth in the Linux sector. Red Hat announced that Red Hat Enterprise Linux subscriptions had grown by over 10 per cent in the quarter to August.

- As ever, www.groklaw.net has more info on these developments than most people have the time to read.

NEWSBYTES

■ The **GNOME project** has set a cash bounty on a number of features that are deemed essential for the future success of the desktop environment, especially in the area of application and cross-platform integration. The awards range from a few dollars – for creating and formatting a range of national and religious calendars for *Evolution* – to US\$2,500 – for integrating the buddy lists in *GAIM* and *Evolution's* address book. The idea, organisers says, is to “help improve the level of integration between some of the core components of the Linux desktop” and, more specifically, “improve the experience of collaboration in the desktop environment” www.gnome.org/bounties/



■ A new version of **Libranet**, which adds a usable, ‘mostly automatic’ installation procedure to Debian, has hit the mirrors. Libranet 2.8.1 features the latest stable Gnome and KDE desktops, as well as the usual enormous selection of applications. <http://libranet.com>

■ An attempt has been made to insert a root exploit back door into the next generation kernel.

An unknown person, in the guise of kernel hacker David Miller, ‘contributed’ the exploit to the file `kernel/exit.c` and changed the user ID of the process to root. The problem was discovered by the Bitkeeper to CVS export scripts and didn’t get into the core *Bitkeeper* repository. Larry McVoy of Bitmover, the developer of *Bitkeeper*, said that though the event was significant, it didn’t really pose a danger to users. “Even with the delay, the problem was identified and corrected in less than 24 hours. That doesn’t leave a lot of time to have the problem get into the real release tree.”

■ The **Mozilla Foundation** has released an alpha version of the latest *Mozilla* browser suite. Many of the big additions are in the mail client, including *vCal* support, a command for removing mail from a server after a specific time period and new signature options. www.mozilla.org

■ An Open Source alternative to Linuxant’s *DriverLoader* (See *News*, last issue) has been established. *NdisWrapper* has been tested with a few WLAN cards, but is still at the early stages of development. <http://ndiswrapper.sourceforge.net>



MarbleBlast is included with the **Wine Rack** package.

BUNDLING LINUX APPLICATIONS

SUSE turns Wine seller

SUSE Linux has launched a new boxed Windows compatibility suite called *Wine Rack*. Pairing Codeweavers’ two core products, *CrossOver Office* and *CrossOver Plugin*, with Transgaming’s *WineX*, the package is available for a wallet-friendly US\$39.99. The box also includes the game *MarbleBlast* from GarageGames. SUSE claims this is a saving of \$117 (or 70 per cent) of the individual purchase prices.

SUSE spokesman Holger Dyroff said the new launch would help break the barrier to adoption for many potential

users. “Customers can opt for Linux, knowing that they can take advantage of both Open Source *and* proprietary applications.” It was, he said, another demonstration of the company’s “commitment to provide the most comprehensive Linux offering for the consumer market.”

Wine Rack allows SUSE Linux users to install and access applications such as *Adobe Photoshop*, *Office XP*, *Quicken* and *Internet Explorer*, as well as view QuickTime files and play games including *Counter Strike*, *Warcraft III* and *EverQuest*.

CONFERENCE

Desktop Linux Consortium

The **DLC** inaugurated its annual

Desktop Linux Conference at the Boston University Corporate Education Center, Massachusetts with a host of speakers including Bruce Perens, Ximian’s Nat Friedman, Jeremy White of Codeweavers and Jim Curtin, CEO of NeTraverse. The main focus was the issue of Windows migration and the establishment of an enterprise-centric Linux distribution.

Perens suggested in his keynote address that the best way to drive adoption of Linux in business circles would be the creation of a user-supported distro tailored to the needs of enterprise. Companies would acquire this User Linux, he said, by ‘paying for software engineering’ rather than on a per-seat basis. The core of User Linux, which would be supported by multiple vendors who would add their own bells and whistles aimed at specific markets, would be Debian; with the non-free elements such as 3D drivers included.



Perens predicted IBM buying Novell and Linux on 30% of enterprise desktops within two years.

Nat Friedman hinted that Novell may sanction the porting of *Evolution* to Windows as an aid to migration. Novell, he said, believes that the client software should remain free, but that it intends to make money on resource management and server software. To that end, Ximian is currently working on a resource management system for *Evolution* and the whole GNOME desktop. It should see the light of day in 2004.

David Cartwright

An IT consultant since the phrase was respectable, David specialises in Linux systems and solutions.



COMMENT

The Grid’s in fashion

“The concept of ‘grid computing’ has become trendy in the last few weeks. Although not trivial to implement technology-wise, the idea is simple: instead of running a program on a single computer, you spread it across a number of separate machines within your organisation in order to run the program faster. It’s particularly handy if you have machines that are busy some of the time and idle the rest (eg accounts systems will probably be lightly loaded for 28 days of the month but will be hammered at months’-end).

Vendors such as IBM are promoting the view that in future, we won’t have “personal” computers, or application-specific ones. Instead, we’ll just have all the machines in the company configured to be part of a single ‘virtual computer’. Regardless of architecture, they’ll all run a Java Virtual Machine, and each time we want to run something, we’ll drop the application and the data onto some clever little scheduling software and it’ll go away and figure out how to chuck the job around the network in order to get it processed most effectively.

This environment is blatantly suited to Linux, since at the lowest level all you need is a nice, slimline operating system kernel to sit your JVM on – without the clutter of built-in Web browsers, OS-specific runtime libraries and all that stuff. If, as seems likely, organisations start to implement grid computing, it’ll be interesting to see whether Microsoft starts to feel threatened, and what sophistry it spins to try to persuade the world that its OS offering is the best thing to run your grid on.

Mailserver

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

★ Letter of the month

This month's winner receives a **256MB Crucial GIZMO USB storage device**

Linux saves school £25000

I am contacting you from a school using Linux for over 90 per cent of our IT requirements. We are an independent school, and as such have not gone the *SIMS* route for school management. We do however have four Windows PCs used for administration, as we have to continue using MS-based accounts/payroll packages, but that is all.

Our server (including Samba for the MS machines)/firewall configuration is all Linux with 99.99% up time for all users. Our class teachers/other staff all use Linux (Red Hat) workstations. Classroom/school management information utilises MySQL and *Rekall* front-end, which is similar to *Access/Paradox* and is cross-

platform/cross server. This has all been developed on a part-time basis with no training or outside help except what is available through the many excellent Linux publications and support websites.

OpenOffice.org is our standard cross-platform office application.

Furthermore, work is under way to use *Rekall* for the teaching of databases to year 6 and up.

Estimated total saving by having an absolute minimum of MS licensing costs over the last 2 years, for a school of 175 pupils: £25,000.

For a secondary school, Linux can provide all that is required. It is unfortunate that so many maintained schools in this country are tied into Microsoft products, but it is not the fault of Linux and Open Source that the school in

your magazine article is unable to make use of all the benefits of the Linux OS.

Schools and other organisations should be encouraged to think very hard about using any applications that are not compliant with modern standards or cross-platform for obvious reasons. The DfES and LEAs ought to have departments generating standard OS applications that can be used throughout educational establishments instead of relying on expensive licensing. Government establishments are very backward in this respect witness the number of *MS Word* downloads that you get from their websites. (Have they not thought of using PDFs?)

I would suggest that the writer of your letter sticks with it if he really

wants to save money and get a better system at the same time.

Phil Marshall, *Managing Director, Beechwood Schools Ltd*

Thanks very much for sharing your experiences with us. I think the general point is that schools get no support from central government for making such decisions on IT, which is a shame because this is an area where a great deal of money can be saved and spent on the things schools should actually be spending their budgets on.

For your endeavours, we are awarding you the Star Prize this month: a handy USB storage device. Oh, and you'll find that *Rekall* is on this month's coverdiscs.



Schools IT newsgroup

After reading a letter about a negative experience in education (LXF47), I thought I would reply with my experiences of Linux infiltrating education.

My first topic is concerning a student project that is going on in a colleague's school where they intend to undercut *OpenOffice.org* distributors: and all in a good cause. You can read an article from the local paper here: www.thekjs.essex.sch.uk/yates/images/pic.jpg

For every two *OpenOffice.org* CDs they sell, they will donate one copy to Malaysia to help improve the kids' education over there. They only cost £2 each and p&p is the



cost of a first-class stamp for 1-5 CDs, 2xFCS for 6-10 CDs and then waived for 10+. The CDs contain Windows, Linux and Solaris versions of *OpenOffice.org* 1.1 and there is spare capacity on the CD for extra goodies of your choice. To order CDs or submit questions, the teacher in charge can be contacted at pburgess@wilmslowhigh.cheshire.sch.uk

Secondly, my personal experience of Linux in schools is very different to that expressed in the letter. I have found that there is a lot of Linux activity going on in UK schools at the moment. A large proportion of schools now use Linux servers for a multitude of tasks, and increasingly Linux workstations are being used as thin

clients. As the applications increase for Linux, then you will see more and more schools adopting Linux for the desktop. This is already starting to happen significantly in further and higher education establishments, but is not unheard of in secondary schools. To discuss Linux issues with teachers, your readers could use the newsgroup uk.education.schools-it, and SUSE also has a schools mailing list.

I'm sure a lot of your readers have an interest in how Linux is used in education and I hope to read articles on this matter in your magazine in the future. Perhaps a like-for-like comparison of Windows and Linux applications that are specifically used in education (both client and server) would be an interesting topic and may encourage others to consider the advantages of migration.

Darren Smith, *Egglecliffe School, Stockton-on-Tees*

Thanks very much for letting us know about the *OpenOffice.org* distribution scheme, and for your own thoughts on Linux penetration in schools. We are always very happy to hear from readers on how Linux can be used in the education sector. If you are interested in implementing Linux in schools, further or higher education, why not visit our forums to discuss the issues you face with other Linux users?

Higher than the Sun

In his column in *LXF48*, if David had got hold of a proper distribution of Solaris (say Solaris 8 12/02) then not only would he have had a set of bootable CDs but he would also have found a very useful CD entitled *Bonus Software*. This CD contains an assortment of precompiled freeware software ranging from X window managers, editors, network tools and lo! A compiler (GCC).

As for the comment about the lack of `make`, take a look at `/usr/ccs/bin`. Of course, if during the install he chose not to load the development package, then these utilities would be missing. But it's not necessary to resort to pulling the stuff of the net, simply load the packages off the CDs. Moreover, there are numerous sites out there that provide access to already compiled packages. One such site is



There's plenty of Solaris/SunOS resources available for free on the Internet if you know exactly where to look.

www.sunfreeware.com. I too can remember the days when SunOS 3.2 was all the rage; gcc was really quite useful for compiling X11Rx. It's really quite amusing that the Linux community thinks it has invented everything... And yes, I do run Linux on one or two PCs plus Solaris on x86/Sparc.

Neil Davis, *via email*

David Cartwright writes: Sounds like the company whose machine I was setting up had secreted the 'bonus' CDs somewhere, because the version of Solaris 8 they'd given me was indeed a "proper" boxed distribution, albeit not a sealed one, and it certainly didn't have any bonus CDs. So it sounds like I've maligned Sun on that front – sorry chaps.

With regard to the precompiled binaries site you mentioned – you're absolutely right, and in fact it was one of the mirrors of Steve Christensen's www.sunfreeware.com that I found myself downloading binaries from. I've used this loads of times over the years, and it never fails to amaze me that among the junk one finds on the Internet one comes across these mines of usefulness.

Spam, spam, spam

Thanks for the very interesting article on spam. It covered a lot of ground, and explained how spammers get money out of the whole process, which had been a mystery to me before now.

The fact that the government of the UK is considering spam as a legitimate form of marketing would indicate that MPs haven't quite cottoned on to the realities of some aspects of the Internet. They are trying to promote e-government, yet not even all of them have email, so how can they comment on spam?

I read recently of an interesting retaliation on a well-known spammer in America, Alan Ralsky ([www.counterpane.com/crypto-](http://www.counterpane.com/crypto-gram-0304.html#1)

[gram-0304.html#1](http://www.apig.org.uk/)). Strangely, he made public his home address and a group of his spam victims started registering him online to any and all mail-shots and catalogues they could find on the web. Within days this physical spam started arriving at his house, and became so much that he couldn't get out of his front door.

Rob Purvis, *Exeter*

You're quite right about the UK government appearing misinformed/behind the times/dangerously ignorant about even the basics of electronic communication and the Internet. If you want a cheap chuckle at their expense, have a read through some of the information that is provided at www.apig.org.uk/ – not a bacon-farmer's or Police site as you might expect from the URL, but the UK Parliament's *All Party Internet Group*. The fact that Chairman of APiG, Derek Wyatt MP, stated in public that the spam problem could be eliminated by appending postcodes to people's email addresses just beggars belief.

There are some good points made in isolation in some of the materials downloadable, but altogether a lack of coherent thinking demonstrated by those in positions of political power about electronic media in general makes us at *LXF* highly concerned for the future. From recent fiascos concerning e-voting in both this country and the USA, to the great money-gobbling NHS IT debacle (see www.theregister.co.uk/content/7/34276.html), it's enough to make even the most fervent advocates of personal freedom and properly accountable democracy wonder whether the world would actually be a better place if it was completely governed by big corporations...

Welcome back!

In the field of software time is measured in version numbers. Soon after the arrival of SUSE Linux 7.2, I gave up trying to use Linux as my

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: lx.f.letters@futurenet.co.uk

operating system of choice. I work from home and rely heavily on my computer for my work. Consequently, I did not have the time to make Linux work for me. With the arrival of Windows XP I thought my flirtation with 'the penguin' was over. That is, until I found myself with a spare moment and installed Red Hat 9.0. The installation was easy, every piece of hardware was recognised correctly and worked first time, and I found myself able to concentrate on exploring the goodies on my hard drive. The moral of this story?

1 My system is usable (and reasonably reliable) for my work, therefore I can now explore Linux more deeply. With this in mind could you recommend a good book for someone who has been using computers for a while and who knows his way around a 'Repeat ... Until' loop but feels awed by the apparent complexity of Linux?

2 I have now renewed my subscription to *LinuxFormat*.

3 Red Hat lost out to SUSE on my computer because of the apparent 'withdrawal' of the company's commitment to the desktop computer, though I would be interested in trying Fedora if you ever put it on the cover disk.

4 Now that my PC can do virtually everything I want I find myself thinking more about the philosophy behind Open Source software. It's good to be back!

Keith Hulse

Glad to have you back. We have reviewed a number of books on Red Hat recently – see last issue's book reviews, or type "Hoyt Unleashed" into Google. Fedora is on this month's coverdisc, as we hope you've noticed! >>

READER TIPS

Reverse lookups

Regarding your answer *Reverse lookups* (p98, LXF46) regarding reverse lookups when you only have part of a class C network range. Partial class delegation is possible: see *rfc-2317*.

Alternatively, have a look at www.acmebw.com/askmrdns/arc_hive.php?question=7

Or Google gives pretty good results: www.google.com/search?hl=en&ie=UTF-8&oe=UTF-8&q=howto+reverse+classless+dns&btnG=Google+Search

J David Rye, *via email*

In LXF46 there was a question about reverse lookups on partial class-C ranges. The answer was 'not possible'.

Check out *rfc-2317* ([www.faqs.org/ftp/rfc/pdf/rfc2317.7.txt.pdf](http://www.faqs.org/ftp/rfc/pdf/rfc2317.txt.pdf)) for a description of the so-called 'CNAME hack' for one way to get around the problem. It needs the co-operation of the ISP to set it up, but once set up the DNS sub-range delegate can control the reverse lookups.

John Walsh, *via email*



Thanks to *OpenOffice.org*, Linux's support for non-Roman characters and right-to-left languages beats any other OS you can think of hands down.

Speak in tongues

I was pleasantly surprised that I am not the only one seeking a way to mix Roman alphabet with Japanese characters in one document (*Turning Japanese, Answers LXF47*), in my case *OpenOffice.org*. Although the use of Asian Characters (and other scripts such as Arabic, Thai, etc) is touted as A Big Selling Point for *OpenOffice.org*, as good as no information is available on how to get these into actual documents using a Dutch or English version!

I've been able to enter Japanese script in *OOo* (*Writer, Calc* and *Impress*), with the aid of *kinput2* and *Canna*, which is not difficult to install with SUSE 8.2. I didn't realise I had everything necessary installed until I tried to use it! It seems to be enough to mark *Canna* and *kinput2* in *YaST2*, and fill in the options for Asian languages in *OOo* (Tools>Configure), to integrate these into the boot sequence and ready for use. I am not sure which fonts are necessary, though. Mandrake 9.1 and Red Hat

9 aren't as Japanese-friendly, as far as I'm aware. Fellow Nipponophiles may find this pair of links useful :

<http://www.suse.de/~mfabian/> Mike Fabian's internationalisation site, with lots of info on Japanese and a bit on Chinese and Korean too. Useful for other distros too. Also read *CJKV Information Processing* by Ken Lunde (ISBN 1-5659-2224-7, O'Reilly). <http://www.nec.co.jp/canna/index.html> (Japanese text only !)

Dirk de Hoon, *Belgium*

MDK hardware

Thought I would let you know some good news about hardware in Mandrake 9.2. I got an old Mustek Scanner - 1200UB. This one was discontinued because of a hardware bug, but SANE talks to it fine right up to the full definition of 1200dpi, the official Windows driver won't clock it, though Windows 98SE OKs it with its own diagnostics!


Also I bought a cheap digital camera for £13 from the cheap store. *GPhoto2* recognises it as an Argus 1510 and downloads from it without fuss, bearing in mind the cheap manufacture !

Bryan, *via email*

Helpdex

shane_collinge@yahoo.com





See also www.kyfeithu.co.uk and www.gwelywiwr.org for Linux in Welsh.

CymruLinux

I particularly enjoyed the Mandrake Linux 9.2 distro on LXF47. The reason for this is Mandrake's particular strength in its language support. Included in this version is a full Welsh version of its own software and substantial amount of the GNOME desktop. Add to this the Welsh language pack for *Mozilla* and it's fast becoming *the* distro for Welsh language speakers.

To access this feature either chose Europe>Cymraeg at the installation stage or change the locale to Cymraeg at the desktop.

Interested? There's more info available (mainly in Welsh, sorry!) at www.meddal.org.uk website on a variety of Linux and non-Linux software in Welsh.

Rhoslyn Prys, *via email*

ISO Image collapse

Having retired, I have time to start dipping a toe in the Linux water, and I decided to install Mandrake 9.2 on my older box (Pentium II 300MHz, 512MB RAM, 6GB, 5GB & 40GB HDD, DVD-ROM & combo DVD-ROM/CD rewriter). Having wiped all the hard disks, I tried booting from the DVD (the DVD-ROM drive is the secondary master) by setting up the BIOS to boot as first choice from the CD, but it wasn't having any of it, so I decided



to drop back to burning CDs on my Windows machine. I kicked off the *winmkiso* script and let it chug away; fortunately as each ISO image was completed I moved it out of C:\ to another directory, because I found that the fourth ISO image was created as Mandrake-9.2-Install-CD3.i586.iso – oops! I renamed the fourth ISO image to Mandrake-9.2-Install-CD4.i586.iso, and moved it into the same directory as the others.

You may like to warn readers to watch out for this, as otherwise they might be rather puzzled about the four images collapsing to three. Ian Park, *via email*

SpamAssassin

Your FAQ on *SpamAssassin* in LXF47 was interesting, but clusters? *procmail*? *CPAN*? Geez, that'll send the newbies running back to their Windows boxes for sure! For us mere mortals with simple needs it can be a lot less daunting, especially if you use *KMail* with a dial-up connection. Dr. Tony Young's letter in *Reader Tips: Kill Spam Dead* was along the right lines for your average desktop user, and using his

KMail 'filters' method together with *SpamAssassin* really does kill spam dead – well, 99% of the time.

SpamAssassin (SA) is included with some recent distros. It is easily integrated with *KMail* and instructions for a SUSE 8.2 Pro system (SA not included in SUSE 8.2 personal) are included below and can be adapted for other distros.

[NOTE: This information is adapted from instructions found at: <http://kmail.kde.org/tools.html>.]

Checking your email is quicker if you run *SpamAssassin* as a daemon in the background. If you don't want to do this, then skip item 2 below and replace **spamc** with **spamassassin** in item 4i).

1 Install *SpamAssassin* with *YaST2*> Install And Remove Software (this will include several dependent Perl packages).

2 Staying in *YaST2*, go to System> Runlevel Editor. Click on the button 'Runlevel Properties' and scroll down until you find the entry 'spamd' and highlight it. Check the '5' box so that *SpamAssassin* (*spamd*) will run as a daemon in the background during any normal GUI login session. Start *spamd* running with the 'Start/Stop/Refresh' button. Click on the 'Finish' button and exit *YaST2*.

3 In *KMail*, add a new folder called 'Spam' (Folder>New Folder).

4 In *KMail* go to Settings>Configure Filters. Next, we need to add two filters here:

i) Filter Criteria: Match all of the following:

<size> is less than 250000

Filter Actions: pipe through>spamc

Advanced Options: DESELECT – If

this filter matches, stop processing here.

The above filter size is necessary as *SpamAssassin* apparently throws a wobbly on emails bigger than this, but any files smaller than this will be sent through *SpamAssassin* and if identified as spam will be tagged as such. The next filter tells *KMail* what it should be doing with the tagged emails.

ii) Filter Criteria: Match all of the following:

X-Spam-Flag equals YES

Filter Actions: move to folder>Spam

Ensure the size filter is at the top of the list, and that's it. Next time you check your mail, you should see all that horrible spam being piped into your new Spam folder; or if you feel confident, straight into your trash folder. However, in practice I do find the odd non-spam email going into the spam folder and the odd bit of spam getting into my Inbox. This is reduced if you move any messages that have slipped through into their correct folders and occasionally run the following commands from a terminal.

For NONSPAM

```
sa-learn --nospam --dir
~/Mail/inbox/cur"
```

For SPAM:

```
sa-learn --spam --dir
~/Mail/Spam/cur"
```

These commands 'teach' *SpamAssassin* so it may better identify spam (and non-spam) in the future.

Hope this is of help!

Martin Collins, *via email* [LXF](mailto:mf@lxf.co.uk)



If you're seeing too much spam in your mailbox, *SpamAssassin* can be integrated with nearly all mail servers.

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...

HP DW530 and xw4100 desktop workstations >>

Sleek, silent and super-fast – has HP's dedication to its Linux partners made this the definitive workstation range? **p18**

Omnis Studio 3.3

Rapid Application Development environment produces usable results in 30 minutes! **p20**

Red Hat Enterprise Linux 3

Over 100 new features and 350 general enhancements to AS, ES and WS since 2.1 **p22**

Maya 5

3D content creation tool gets the vital functionality that it has been lacking until now **p24**

VIA EPIA CL10000

Good things come in small



packages: this mini-ITX board is no exception to the rule **p27**

Smoothwall Corp. Guardian 3.0

If you can't trust your users, this will defend your network by

preventing them from accessing unsavoury web content **p28**

Book reviews

Masters Of Doom, Beyond Fear, OpenGL Game Programming, Text Processing in Python **p30**

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
APACHE	1.22
MYSQL	1.11
COMPILE	0.96
OGGENC	1.71
OVERALL	1.17

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

LINUX DESKTOP WORKSTATIONS

HP's little wonders

Sleek, silent, and super-fast – has HP cracked the Linux desktop market?

Paul Hudson reviews two machines that aim to please...

BUYER INFO

Low-end workstations somewhere between the (discontinued) Evesham Quest LM machine reviewed in *LXF43* and the Armari Pro3D workstation reviewed in *LXF46*.

HP WORKSTATIONS, d530 AND xw4100**■ SUPPLIER** HP

■ PRICES d530: £790+VAT, xw4100: £1405+VAT

■ WEB www.hp.com

As computer manufacturers go, HP has an unparalleled history of quality in both the design of its machines and the quality of manufacturing that characterises its hardware across the board. With the Linux workstation market on the apparent verge of reaching critical mass, several vendors are rushing to get into this lucrative market. HP, never one to rush, has clearly been putting a lot of thought to how the company can take its extensive market knowledge and apply it to Linux.

Having reviewed HP's Itanium 2 server in the last issue of *LXF*, HP's high-performance computing expertise is certainly well-proven. For workstations though, we've only seen the various press releases about partnerships with Red Hat, SUSE, and Mandrake fly by in the past, and only now is HP sending the first hardware

out onto the market running Linux. Fortunately, we've got two machines in for review...

Big screen dolls

Taking these two HP workstations out of their boxes, you're immediately hit by the aesthetic lengths HP has clearly plumbed to make these machines look and feel worth their money. Both machines come in a two-tone silver and black colour scheme, with smooth curves and a fairly uncluttered front plate.

Specification-wise, the xw4100 comes with a 2.6GHz P4, 512MB RAM, 80GB hard drive, a combined CD-RW/DVD, and Red Hat 7.3 pre-installed. The D530 came with a 2.8GHz P4, 256MB of RAM, a 40GB hard drive, a CD-ROM drive, and Mandrake 9.1 installed, curiously branded Mandrake 9.1 Light – presumably that means it only has 3 CDs, as opposed to the full ProSuite edition. On the upside, the D530 has a gigabit network card built in, and so has a convincing edge on the network.



HP d530 A good option for customers disappointed by Evesham withdrawing from the UK Linux PC market.

respect, it's best to go for the xw4100, despite the fact that it has a slower CPU than the D530, because it has twice the RAM.

Intergalactic

Quality components plus a household name brand should mean performance, performance, performance, so we were quite eager to put these machines to the test. Because these machines don't have SATA drives or Xeon CPUs, we weren't expecting groundbreaking results, and so were quite amazed as the results came in. The benchmarks are fairly close together, with fairly predictable score differences between the two.

With a hard drive score of 3.28, the xw4100 is a screamingly fast machine, particularly when you take into account the fact that it uses ATA technology as opposed to SATA or even SCSI. Given that a score of one means that machine equals the SCSI drive of our yardstick machine, it's quite a feat for an IDE machine to get such a high score, and reflects very well on HP. This high performance continues in both the *Apache* test, where a score of 3 seems to be about average with the current generation of machines, and also kernel compilation, where 1.82 is quite respectable also.

Perhaps the high point of these machines is their MySQL scores – they are both consistently high, particularly for a machine of this price. While workstations aren't often used for databases, they certainly do have to put up with quite a lot of random access data retrieval, so this score is quite important.

Overall scores of 3.14 and 3.04 are very strong scores, and compares very

As workstations go, the xw4100 comes in quite a way ahead of the d530 as far as usability goes – it's got a little less power, yes, but it more than makes up for that by having an nVIDIA Quadro4 980XGL graphics card with 128MB RAM to make it into a graphics powerhouse. Of course, the fact that it's got a bundled CD-RW/DVD drive is also good, and the combination together makes this a strong multimedia contender. Comparatively, the d530 only has on-board graphics – more than enough for all 2D needs, but simply not up to scratch if you want any sort of 3D ability.

As you can see, both of these are fairly high-powered boxes with enough RAM and hard disk space to be suitable for most uses. They clearly aren't in the super high-end category – 256MB of RAM is likely to run dry if you need a data cache in your applications or engage in complex work such as 3D modelling. In this



Connection ports should be available on the front of boxes as standard.

well with machines two or three times the prices seen here. While it's possible (and indeed probable) that we'd see even further improved performance if the hard drives were swapped for SATA or even SCSI drives, it would also ramp the price up considerably, therefore taking away the value aspect that is the foremost advantage of these machines.

Sit on my karma

Both of these machines are extraordinarily quiet, which is quite a surprise as they each have two large fans at the back. If you're any further than six inches away from the rear fans, it's very hard to pick out any sound at all, even when the machines are thrashing the hard disks. We haven't the first clue how this is possible, but the adage about the mouths of gift horses comes to mind!

In line with HP's near-legendary design expertise, both machines come with two USB ports on the front, as

well as headphone and microphone ports. Around the back there are four(!) more USB ports, as well as the usual array of ports and slots, so there's certainly enough connectivity for most people here.

As well as the machine, there's actually quite a lot of other kit in the box, as we've come to expect from HP machines. Apart from space power and network cables, there's also a set of Mandrake 9.1 CDs, a 10-page Mandrake mini-guide, a Quick Setup leaflet, and a Getting Started book – more than enough documentation to get users up and running. Added to that, there's a matching silver/black keyboard and mouse, just to finish off the sleek and modern look.

Note that neither of these machines come with any monitor included, so you'll need to provide one yourself. Nevertheless, as far as added-value periphery goes, these machines definitely come top of the table.



HP xw4100 An attractive budget desktop solution, despite the Windows logo!

Happy nation

Despite being from one of the most well-known computing brands on the planet, and also despite having high-power components designed to satisfy even demanding users, these machines both still manage to come in at affordable prices. Few could argue that £790 is expensive for the d530, particularly when backed up by HP's warranty and support. The added cost of the xw4100 (£1405) is largely justified by the very expensive graphics card (it's certified for use with a variety of graphics applications), but unless you really intend on making the most of the graphics card bundled with it you're probably best off requesting a machine without it and saving yourself the extra £400 cost.

On the performance front, both machines put forward the kind of convincing performance you'd expect from such a calibre brand. We were particularly impressed with both the hard drive tests and the MySQL tests, as the scores are a great deal higher than we, or anyone to be honest, would expect from this configuration. HP have made quite a few nips and tucks to the specs here to try to keep costs low – many companies opt for the very latest Intel CPUs and SATA hard drives; and yes, these latest components do yield improved performance, but they do so at a high cost. HP, on the other hand, has used CPUs that are a bit older, and passed the savings onto the customer. If you want to request the latest and greatest, we're sure that HP will happily comply for a suitable fee – you get what you pay for.

As there's a choice of Linux distros, we'd recommend you go for the Mandrake option – the Red Hat 7.1 release shipped was quite old, and Mandrake 9.1 continues to be the easiest-to-use distro available. HP does have a long-standing deal with SUSE, so you may well have a third choice for distro if you ask, and it's not likely to nudge the cost up by much.

Thanks largely to its size and commercial weight, it's not surprising that HP can offer a warranty any business user would be happy to rely on – it's a three-year, next business day, on-site hardware and labour warranty, which is pretty much as good as it gets. We're not sure how HP can profitably manage to squeeze this warranty into the low prices here, but it will definitely please purchasers looking for enterprise-level reliability.

Overall, these are two cracking machines that provide above-average performance across the board, often even very much above average. The only machines we've seen for a while that cost-effectively perform this well have been Opteron-based, so it's good to see that HP is managing to keep competition alive with these boxes.

At this price most people wouldn't think the machines would come with a choice of distro, or even such a comprehensive warranty package, but the clear fact is that they do, and we think you'd be crazy not to take advantage of it. In the meantime, would someone at HP please tell us how they managed to get these machines so quiet? [LXF](#)

BENCHMARKS: d530

HD	3.05
APACHE	3.35
MYSQL	3.21
COMPILE	1.82
OGGENC	4.31
OVERALL	3.14

BENCHMARKS: xw4100

HD	3.28
APACHE	2.98
MYSQL	3.16
COMPILE	1.78
OGGENC	4.0
OVERALL	3.04

LINUX FORMAT VERDICT

HP d530:	
FEATURES	8/10
PERFORMANCE	10/10
EASE OF USE	9/10
VALUE FOR MONEY	10/10

A powerful machine with a high overall score, but lacking the DVD drive and decent video card some users may want.

RATING **9/10**



LINUX FORMAT VERDICT

HP xw4100	
FEATURES	10/10
PERFORMANCE	9/10
EASE OF USE	9/10
VALUE FOR MONEY	9/10

Not quite as fast, but still packs a powerful punch and with extra features to boot.

RATING **9/10**



RAPID APPLICATION DEVELOPMENT

Omnis Studio 3.3

Cited by its creators as “one of the most powerful and cost-effective RAD tools”; Jono Bacon knows why...

BUYER INFO

Easy-to-use RAD tool that could use a little final polish – for alternatives, see below and box on opposite page.

- **SUPPLIER** RainingData
- **PRICE** From 250 Euros (Standard edition)
- **WEB** www.omnis.net/

Rapid Application Development (RAD) tools are a common occurrence for Windows, and although we have been treated to the likes of *Kylix*, *JBuilder* (see LXF44) and *MTStudio* (see LXF47) on Linux, it's still true to say that there are relatively few RAD tools for our favourite OS. One of the contenders for the RAD market is RainingData with its *Omnis Studio* 3.3 release.

Before I received my review copy of *Omnis Studio*, I performed my normal background checking about the product, and encountered a quite impressive amount of PR on the RainingData website. Although there was the usual amount of marketing hyperbole, this application was looking distinctly interesting. A few days later I received my review copy and started it up straight away.

Installation

I slotted the disc into the drive, mounted it and proceeded to start the setup program in the linux directory on the disc. Up popped a nice Windows-

style installation routine and I selected New User so it would install the typical tools. A second later I was told that the packages were installed. A bit suspicious of the speed, I looked in the shell window where I started the setup program and it said that it could not read the RPM database – not surprising as I was running Debian. I was advised to use *Alien* to convert the RPM to a DEB, which I did. I then installed the DEB, moved to the `/usr/local/rainingdata/os33` and unpacked a tarball.


I found this whole process a bit of a pain, and RainingData certainly need to work on making non-RPM installations smoother; after all, I would not imagine that creating a DEB would take all that much extra work. Debian is a common distribution and it still surprises me that so many commercial Linux products do not support it.

First Impressions

Due to the slightly unusual nature of installing *Omnis Studio*, I had no icon available on my KDE desktop or menu, so I started the application from the shell. I was instantly thrown back to the prompt and told that a shared library could not be found. I set `LD_LIBRARY_PATH` to the omnis directory and then successfully fired up the application. Given the installation woes I encountered, I had been hoping that things were going to go a bit more smoothly...

On the CD and DVD

In order to give you the opportunity to try out some of the features of *Omnis Studio* for yourself, we have secured a deal from those generous people at RainingData to include the FULL VERSION of *Omnis Studio 3.1* on this month's coverdiscs. Give it a whirl and see what you think!



RAD development

Omnis Studio has a number of dialog windows that represent different parts of the application. When you begin, you create a new library (similar to a project, and stores different parts of your application such as your database, forms etc). As well as the library window, we have the main *Omnis Studio* window with application menus and toolbars and a components dialog. This components area has a number of different types of component that can be used within your application. To use a component, you can drag it over to the library, and then a properties window will pop up to allow you to configure the component properly. One of the nice features that I liked about *Omnis Studio* is that many of these components have a wizard to help you get it configured properly. This was a nice touch and greatly helped getting some of the more complex components set up easily.

One of the first components to add to a new application is a database widget. *Omnis Studio* supports a number of databases, and lets you

When *Omnis Studio* loaded up, an attractive splash screen introduces you to the application. Not only does this splash screen describe what *Omnis Studio* is, but it also provides a gentle introduction for new users and a full tutorial. This was a good start and there is a definite feeling that each need has been taken account for. There was nothing that I was unsure about when starting to use the application, and I started going through the tutorial.

The tutorial is a fairly expansive introduction to creating a picture database and shows you how to perform all of the steps to create a database, a form to view and add data, and create a remote form to access the database from the Internet using a web browser. The tutorial takes place over a number of steps, and the instructions are clear and simple to follow, without being patronising as so many application tutorials can be. A complete Linux novice might struggle to follow all the instructions, but you won't need a huge amount of experience to follow along. Before long I had my first *Omnis Studio* application ready.



When the application starts, a friendly introduction gets you started.


easily set up the fields and their data types. This component also lets you configure indexes for the database. Another nice feature of *Omnis Studio* is its included database that lets you create an application without having any SQL database knowledge.

With the database complete, you can then add a number of forms that provide an interface for the application. The form component spawns a wizard that lets you set up various parts of the form and allows you to select from a number of form themes to create a design for your application. With the form complete, you can then create a data source (a place where your data is stored), and begin adding data to your application via your form interface.


Up to this point, and even though I had only used a subset of the features in *Omnis Studio*, I was very impressed with the sheer ease-of-use that characterises the application as a whole. At each step of the way the process is intuitive and documented in an exemplary fashion. I was surprised that in such a short series of steps that I was able to create a database-driven application with a form to interact with the data. For simple storage of data, my current application could be enough.

Further features

Omnis Studio has a quite staggering array of features. In addition to its simple interface to creating applications, the environment can be



The environment uses many different dialogs to perform different processes, which looks cluttered at first, but gets instinctive with use.



Extensive help is available, and *Omnis Studio* has a great range of wizards.

COMPETING PRODUCTS

Other RAD programs for Linux

Kylix

This is the main competitor to *Omnis Studio* and gives Linux a Delphi environment with support for database applications.

www.borland.com/

Rekall

Rekall from theKompany.com provides a

very similar feature set to *Omnis Studio* and is also a cross-platform product.

www.thekompany.com/

Forms Developer 6i

This is an Oracle interface development system for creating Oracle databases with forms.

www.oracle.com/

used to create forms easily with a variety of form components. Database administration is easily executed with the SQL access, and you can graphically administer a number of different databases such as MySQL, Oracle, SQL Server, *Informix* and others. Applications can also be scripted using a built-in language with a number of methods and functions to optimise applications, and reports can be generated with data graphed and analysed with the application.

Another interesting feature with *Omnis Studio* is its web plug-in that can be used for accessing *Omnis Studio*-based applications. These applications can run using the server part of the studio, and this provides a more interactive means of deploying your completed application on the web.

Documentation

When I received my review copy, I just got a DVD case with the CD-ROM in it, so I cannot comment on any printed documentation that may be included with the full on-sale edition, but one thing I can comment on is the sheer amount of online documentation included with *Omnis Studio*. There are a number of help topics covering using the environment, using the scripting interface, dealing with databases, deploying your application and other subjects. The documentation seems concise yet expansive, and I was impressed with it.

Combine this help with the wizards, tutorial and introductory guide, and you get a feeling that RainingData has spent a lot of time getting the documentation right for *Omnis Studio*.

Conclusion

At first sight, *Omnis Studio* looks similar to something such as *Microsoft Access*. Although a reasonable comparison to make, *Omnis Studio* offers much more than you would expect, and I only had time to cover a

mere subset of the mass of features in this environment. Installation was a clunky affair, and it is the main downfall of *Omnis Studio*. RainingData needs to focus on some of these installation niggles, and are advised to provide clear instructions about installing on Debian or preferably include DEB files on the CD-ROM – it took me a whole two minutes to create them, and I don't see why RainingData could not manage this.

Other than these small installation problems, *Omnis Studio* is a strong and powerful piece of software. The environment provides what I would consider the essentials in a RAD environment (wizards, form editor, database access etc), and then adds a stack more features and components to ease each area of development. The creation of applications was a simple process, and creating effective real-world functionality was easily achievable within the first 30 minutes of my use of *Omnis Studio*.

If you are looking for RAD environment that works solidly and interoperates well with a web interface to provide access to your data, I would highly recommend *Omnis Studio*. The only real addition that I would like to see would be support for different scripting languages such as PHP and Python, but other than that, *Omnis Studio* is a highly recommended, powerful piece of software. **LXF**

LINUX FORMAT VERDICT

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

With install wrinkles and other annoyances ironed out, and coherent support for a variety of scripting languages, a later iteration could be a future market leader...

RATING **8/10**



3D MODELLING & ANIMATION

Maya 5 Unlimited

Nick Veitch enters another dimension with the latest port of this 3D powerhouse.

BUYER INFO

3D modeller and renderer. Competition on Linux from *Realsoft 3D*, others may follow.

- **DISTRIBUTOR** Alias|Wavefront
- **PRICE** From £1,499
- **WEB** www.alias.com

Alias|Wavefront has a long history of producing top-of-the-range 3D tools, and *Maya* is possibly the best-known 3D software on any platform. Version 4 was the first to be released for the Linux platform, and since *Maya 5* was released for Windows earlier in the year, the Linux version is less behind than it was.

Mental Ray


One of the long-standing criticisms of *Maya* has been that the quality of the render engine. Obviously, this is a fairly key area for a 3D graphics program. Work on resolving these shortcomings resulted in the *mental ray* plugin renderer being released just over a year ago, and a new and updated version of this is now included with *Maya 5*.

In addition to improvements in baking, significant improvements have

also been made to the hardware rendering. On the Linux platform, this includes support for ATI Fire GL and nVIDIA Quattro high end graphics cards (which are about the only ones that have proper Linux drivers anyway). Hardware rendering can speed image creation by a factor of 10 to 20, and although the image quality isn't as high as with software rendering, it is nevertheless more than adequate for visualisations. The increase in range of effects that can be handled with hardware rendering – from particles to bumpmaps and reflections – has also been ramped up sufficiently from version 4 to allow a more realistic approximation of the final scene.

One of the *Maya 5* features enjoyed on other platforms, but not available in Linux is the ability to render scenes in vector format. It's one of the latest fashions, but also has a practical side in that it makes rendering for Flash, SVG and other web vector formats more feasible.

Perhaps one of the more surprising things about the new features is that they have taken this



SYSTEM REQUIREMENTS

PII 300Mhz or better
i86 processor
Ethernet card
256MB disk space, 512MB RAM
OpenGL


long to get in to *Maya*! Extruding along a curve has been possible in just about every other modelling software for years, but this is the first time *Maya* users will have seen it.

Fur goodness sake


Fur may be render-intensive, but it's still the best way to render any number of objects. *Maya* uses a fur system very similar to the Paint effects, where fur filaments are applied perpendicular to the vertices of a mesh. New features in

the latest version include controls to adjust the 'clumping' effect of fur, to reproduce a variety of effect, such as wet or matted fur.

Also from the 'quick and easy' department are changes to the paint effects. These are one of the most



Particles and paint effects combine to build up a complex scene.



The interface looks simple to start with, but it's easy to get lost.




powerful features of *Maya*, where complex shapes can be 'painted' onto surfaces (the 3D equivalent of the paintbrush feature in 2D software). This

naturally lends itself to natural items such as plants, grass, trees and the like, so it's no surprise that these feature strongly in the supplied

models. The software has been rewritten to allow more complicated mesh objects and multiple textures, so it's now possible to paint with complete plants and trees. Adding your own objects to the repertoire is fairly straightforward.



The 'Visor' is a library of almost drag-'n'-drop elements and effects.



Modelling capabilities have been improved since version 4.

IK Blending

Both forward and inverse kinematics are fairly standard features in 3D software. These allow jointed objects and models to be moved intelligently towards goal points. For complicated motion though, strictly animated IK systems can still appear jerky and slightly unreal.

Maya 5 now incorporates a system of blending between animation controls applied directly and those inherited through the IK chain. With careful use, this can be used to modify positions and 'nudge' the IK behaviour back into a more realistic animation than would otherwise be possible.

Another new feature on the animation side is an addition to 'ghosting', which is much like the onion-skinning techniques of traditional 2D animation. This previews the animation for a sequence of frames before and after the current one, and has been extended in the latest version to include the ghosting of connected skeletons or other hierarchical objects.

Another improvement of note for those concerned with reducing complex models, particularly for use in games. Polygon reduction algorithms used to be a very hit and miss affair, as it was easy to throw away important detail. The new system allows high detail areas to be 'painted' out, to avoid such unpleasanties. Another tweak is that entire groups of models can be reduced now, instead of just individual objects.

Maya is probably one of the best-known 3D packages, but it certainly isn't easy to use. It's no surprise that a DVD introduction to the software is included. One of the troublesome things about 3D software in general is that, unlike wordprocessors or even DTP software for example, there is no sense of common heritage or a familiar way of tackling the same problems. If you are new to *Maya*, it's going to take many, many hours to understand it. A good resource is the online tutorials site <http://www.alias.com/eng/community/tutorials/>, which will also give you a good idea of what is possible with this software.

Trial – an error?

Unlike the Windows version, there is no 'free' personal edition available for Linux at the moment. This is a little harsh, as it would at least give potential customers the chance to see whether they could get to grips with the software, and indeed whether it would run on their Linux systems, without having to pay for it. As it is, the *Complete* and *Unlimited* versions of *Maya* are towards the top end of the 3D software price bracket.

If you are an existing user of *Maya* version 4, the choice of whether to upgrade or not is really dependent on whether you find the new features essential. Although this is a full point release, many of the improvements are a little esoteric, apart from the changes to the render engine. **LXF**

LINUX FORMAT VERDICT

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

Maya continues to be a complex but powerful tool for professional use.

RATING 9/10



BUSINESS DISTRO

Red Hat Enterprise Linux AS 3.0

Red Hat pulls back and devotes its resources to the Enterprise, and Jono Bacon takes its flagship high-end product for a ride.

BUYER INFO

Red Hat claims that Over 100 new features and 350 general enhancements have been made to AS, ES and WS variants since 2.1

- **DEVELOPER** Red Hat
- **PRICE** See Pricing box (below)
- **WEB** www.redhat.com/software/rhel

Red Hat is one of the most prolific companies in the Linux industry, and arguably the most visible company to those businesses wanting to explore Linux as a possible solution. Recently Red Hat decided to discontinue its consumer retail version of Red Hat Linux, and is instead concentrating on enterprise products. With this, version 3.0 of Red Hat Enterprise Linux (RHEL) product, is this new focus mean that the company is increasing its already respectable responsibility in the enterprise sphere?

When I first got the task of reviewing RHEL AS 3.0, I went online to get some more information about the product. I was shortly assigned a special serial number in which I would be able to log into the Red Hat Network and access the special support services and product downloads. Entering the code was fine and I was then given access to a number of software channels in which I could download RPMs and ISOs. I spent a little time trying to hunt down the ISOs, but could not find them. I later found out that there would be a delay between submitting the serial number and getting the ISOs due to the system updating itself. This was a little confusing, but they appeared in the software channels about 15 minutes or so later, and I was throttling my broadband downloading the four ISOs. Connection speed was good and I was overall impressed with the Red Hat Network at this early stage.

While downloading, I found an 0800 number (freephone number in the UK) for support enquiries, and



The range of applications with RHEL is certainly exhaustive...

thought I would give them a call to test them out. I spoke to a chap called Rob Nunn on the phone, and I was incredibly impressed with his service. He was polite, knowledgeable and eager to help. This kind of customer service was incredible; I can safely say I have never had customer service as good from any company as I did with Red Hat. We were off to a good start!

Installation

After a few hours and a trip to the cinema later, my four ISOs had downloaded, so I burned them to CD and went to install them on my main machine. This machine consists of an Athlon 750, 368MB RAM, 80GB disk space, Belkin wireless card, 3COM Ethernet card and two sound cards. I slotted the disc in, booted from it and was launched into the graphical installer. The first thing I noticed at this point was the quality of the installer; it was clean, simple and attractive to use, and each step in the installation routine was well documented. The usual steps of selecting your language and configuring a timezone were there, and then I was asked to set up my disk. Partitioning was a particularly

simple routine and can be handled automatically or using the Disk Druid tool – this made dealing with multiple disks an easier affair than some competitors, and a variety of mount points and filesystems were available.

After the basic information was requested in the installation routine, I was then asked what packages I want to be installed. I was able to select a typical system setup with commonly used services, and I selected this

default option to see what would be installed. After this selection, I was able to start package installation, and it continued to ask for each installation disc as it filled my system with packages. When the package selection and installation was complete, I was then asked a few questions and the machine was rebooted for the second stage of the installation. I reset the machine, was presented with a nice graphical *grub* and then booted into Red Hat Enterprise Linux.

The second stage of the installation is where the primary configuration takes place, and once again, an attractive environment is available to perform these steps. The typical areas were configurable at this point – sound cards, network etc. I was impressed to see that both of my sound cards (a Soundblaster Live and M-Audio Delta 44) were found, but when I played the test sounds, I could not hear anything. This came as no surprise, as on this test machine I have been having trouble with both these sound cards – a little tweaking unrelated to the distro put things right soon enough. The next issue for me was getting my wireless card detected. So far *no* Linux distribution has had my Belkin card working yet, and RHEL was no exception. To Red Hat's credit, there are a great deal of wireless cards supported as standard that are not in other

RED HAT PRICING

Annual Subscriptions

	Basic Editions	Standard Editions	Premium Editions
Red Hat Enterprise Linux WS Intel x86 Intel Itanium, AMD64	\$179	\$299 \$792	
Red Hat Enterprise Linux ES Intel x86	\$349	\$799	
Red Hat Enterprise Linux AS Intel x86		\$1499	\$2499
Intel Itanium, AMD64, IBM iSeries & pSeries IBM zSeries & s/390		\$1992 \$15000	\$2988 \$18000

RED HAT ENTERPRISE LINUX PRODUCTS

Scalable products according to customer needs

There are a variety of products that fall under the RHEL moniker, and each of them has been designed and developed to be deployed in particular situations more effectively.

Red Hat Enterprise Linux WS

This version is specifically aimed for smaller scale desktop client systems and workstations. The primary intention for use of this iteration is productivity, software deployment and even being used on a headless machine (*ie* without a dedicated monitor) as part of a cluster.

Red Hat Enterprise Linux ES

This version is designed for mid-range systems and networks that require typical services such as file, print, mail, web and other features.

Red Hat Enterprise Linux AS

This is the high-end mission critical version of RHEL, and this version is designed to be used on systems with a lot of processors, a wealth of RAM and critical, very low failure rate requirements. This edition is of course backed up with extensive support.

distributions. Other than this, my Ethernet card was found fine, and I finished the installation and started up RHEL in its normal state.


In use

I found RHEL to be a generally good system. I was presented with a graphical login and logged in fine, and Bluecurve themes in GNOME and KDE look great. It really does look like both GNOME and KDE applications look the same, and this has a great feeling of consistency that is often lost on Linux systems; although some apps did not use the theme, which was a little flaw in the otherwise stylistically coherent system.


The strength of RHEL is clearly its configuration tools and the stability of the software included. Practically everything I needed to do had some form of graphical tool for it. The *Kickstart Configurator* is good part of

RHEL and I was able to configure the major areas of the actual OS. In addition to the graphical tools for configuration, the command line interface and manual configuration files are all sensible, and it is good to know that both configuration types are easily accessible in RHEL – Red Hat doesn't try to hide the technology from you, it just provides different methods of accessing it.

Package management works well in RHEL and there is support for installing errata/packages via the Red Hat Network, and installing RPM files (and their dependencies). There are some quite clever features within this package management system, and I found it to work flawlessly. Many Linux users (including myself) have slated Red Hat and others in the past for their use of RPM over APT, but this really shows that RPM files can be



... And configuration with the *Kickstart* tool is straightforward.



Package management in RHEL is a breeze, even for APT acolytes!

used well. One disclaimer though, this was tested with Red Hat official RPMs, so there may be issues with third party RPMs in terms of dependencies.

A criticism of earlier versions of RHEL was that the package choice was seriously out of date. While this does lead to greater stability, it can be frustrating not to be able to run the latest technology straight off. Thankfully this has been largely rectified, and this version even includes the only recently released Samba 3.

Documentation on RHEL is quite good, and there are a large amount of docs on the Red Hat Network (there is even a Documentation ISO CD). The documentation is well-written, simple to navigate and timely. This is often an area in which the material can be either difficult to understand, badly translated or simply not available. It is good to see that Red Hat has made a good effort in this area.

The only area that I could not really test with this review was RHEL AS running on multi-processor machines and machines with a huge amount of RAM due to lack of such hardware to test it on, but on a fairly normal desktop machine, performance was good. I know Red Hat have a number of kernels included with RHEL (I noticed a special Athlon kernel being installed), and I did notice a better than normal performance level.

Conclusion

I am impressed. I used to be a Red Hat user, but I ditched Red Hat in preference of Debian several years back, and although I give the latest RH

offering a run for its money every so often, I was genuinely impressed with RHEL. The issue that most impresses me is that Red Hat has gone for the whole package here – not only have the developers worked hard on a good software application, but the installation, documentation and technical support are all there and of a high quality to support the product. This is exactly the kind of support that an enterprise needs to consider a Linux solution, and I am glad that Red Hat has invested well in its product and service line.

If you are working at an enterprise and need a solid Linux enterprise solution, I can recommend RHEL. Although not directly comparable to the alternative enterprise products that SUSE and others offer, I am certainly impressed with this. The only recommendation that I give to Red Hat is the same I give to all Linux distributions – get your hardware support working a bit better!

Other than this small gripe, I would be happy using this for my mission-critical computing. [LXF](#)

LINUX FORMAT VERDICT

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

Evolution in perhaps a Darwinian sense. You'd have to be a real Red Hat devotee to spot some of vaunted improvements behind the general sense of smoothness.

RATING 9/10



MINI-ITX MOTHERBOARD

VIA EPIA CL10000

Nick Veitch puts the small but perfectly performing CL10000 under the microscope. Almost literally.


BUYER INFO

In terms of small form factor fully fledged computers, there is little to touch this, save other EPIA models

- **MANUFACTURER** Via
- **PRICE** around £130
- **WEB** www.viamainboards.com

While obviously aimed at a different market to some of the previous EPIA releases, the CL does also improve on them in small ways. One of the big steps forward is the processor. The model reviewed here ships with the newer Nehemiah 1 GHz processor (as used in the

A single PCI slot is available for expansion, but there aren't many things left to add!



A large but efficient heatsink covers the VIA chipset.

m10000). This is not merely a slight speed increase, but architecturally different. The Nehemiah core includes full-speed FPU and also adds the Intel SSE instructions, which puts it on more of a par with the P4 and Athlon XP processors.

In terms of graphics, the CL10000 uses the same CLE266 chipset for the northbridge. As well as acting as a general controller, this includes the hardware MPEG2 decoding also found in the M series of EPIA boards. However, this is not a truly multimedia implementation, as the CL board lacks the 5.1 sound of the M series board.

While we are comparing the board to its close relative, we should also

Dual LAN ports at the rear make this a dream for networking.



The C3 1GHz processor lives here: it's pretty quiet too!

point out that the CL doesn't include the extra video out connectors found on the M series. These have been replaced with a further 2 accessible USB ports and an extra RJ45 LAN socket, as shown below left.

This reveals the true aim of the board – to be a highly flexible, embeddable network device – either as a mini-server or as a simple black box firewall or bridge. With dual LAN, you can create an effective network gateway, or they could also be used to provide a degree of failover for network availability.

One of the problems of loading more power into a smaller space is one of heat-dissipation, but the design of the chipset has taken this into account. The processor runs with minimum of cooling required, and though both the north and south components of the chipset come topped with an impressive aluminium heatsink, neither got at all hot to the touch during tests. This should make the board ideal for modders as well as everyone else, because even in the close confines of whatever weird enclosure that enthusiastic case-modders may want to cram this into, the board should be thermally safe.

The Linux factor

The small size doesn't mean a lack of connectivity. As well as the four USB 2.0 ports and Ethernet connectors, the board is brimming with options, including FIR, I2C, LVDS, KBMS, three serial port headers and more. Memory is taken care of with a single DDR266 DIMM socket that will take up to 1GB.

Expansion-wise, there is a single PCI slot, though with this particular board, it is a wonder what you might want to add here.

Previously, the main problems with these boards – from a Linux perspective – has been in terms of driver support. While VIA has gone out of its way to provide drivers as best it can, there has historically been a degree of lag between the hardware being released and being fully supported by Linux, especially as a closed-source (or at least less than fully Open) approach for many elements mean the drivers have to be built for a number of different distro and kernel combinations.

Thankfully, as the CLE266 has been in use already in the M series, and the extra Ethernet device is just a duplicate of the existing one, so the support for these is reasonably mature. The MPEG decoder poses a different type of problem, because obviously end-user applications need to know of its existence to benefit. VIA has teamed up with the Xine developers to create a special version of the software which does use this feature, and it's available on at <http://sourceforge.net/projects/epiafb> for everyone to have a look at.

That said, there is still a fair amount of tweaking to be done in order to get the system tuned well. **LXF**

LINUX FORMAT VERDICT

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

Better overall support from distros and end-user applications are the only improvements you could wish for on this otherwise excellent mini-ITX.

RATING **9/10**



REVIEWS Corporate Guardian 3.0

WEB CONTENT FILTER

Corporate Guardian 3.0

Want some protection from the dark side of the Internet? **Chris Denton** surfs in safety.

BUYER INFO

Defends your network from unsavoury web pages, a direct competitor of *Censormet* and *Websense*.

- **DEVELOPER** SmoothWall
- **PRICE** From £200 for 10 users
- **WEB** www.smoothwall.net/

Although it would be nice to assume that everyone on your network is trustworthy and responsible, in many cases a more cautious approach is the wiser option. It is well known that there is a vast amount of unsavoury material on the Internet, ranging from the mildly distasteful to the downright outrageous. The consequences to users if they stumble across this stuff can be very serious, ranging from viruses and spam to disciplinary action at work and even, in rare cases, criminal proceedings. It may be better to remove the temptation by introducing restrictions as to what kind of web pages can be viewed.

Corporate Guardian 3.0 is a product intended for just that purpose. It filters Internet traffic according to a defined set of rules and anything falling foul of the policies is summarily blocked. It is produced by SmoothWall Ltd and has much in common with its popular line of Firewalls. Like those, *Guardian* is a dedicated Linux distro designed to run on a single-use PC. However, an important difference is that because the act of filtering is rather processor-intensive, *Guardian* requires a fairly modern PC or entry-level server to run at optimum capacity. A 2GHz CPU with half a gigabyte of memory is the ideal specification, although the acceptable minimum is a 500MHz processor with 128MB of RAM.

The installation procedure for *Corporate Guardian 3.0* could hardly

be simpler. It runs from a bootable CD and requires little in the way of interaction, save for the setting up of a few passwords. If all goes well everything is done in under 10 minutes, and this is likely to be the first and last time there's a need to work at the console screen itself. Configuration and maintenance are performed via a browser, and are on the whole straightforward thanks to a first rate web-based administration tool. For extra security this is accessed on a non-standard port, preferably with SSL encryption although an unencrypted version is also available.

Documentation

The documentation is pretty good, with a small printed guide to getting started quickly included with the CD. This is fine as far as it goes, but to set the machine up properly a bit more depth is called for. To address this need, a 155-page administrator's guide is provided electronically. Although it will keep a printer busy should working from the screen be a problem, this PDF is comprehensive in its coverage and well-written enough to be easy for users to follow.

Under the bonnet, *Guardian* is essentially a combination of two popular open source applications, *Squid* and *DansGuardian*. *Squid* provides the proxying so that client PCs are not themselves connected to the Internet. It also uses a cache to satisfy requests locally where possible, thus saving bandwidth. *DansGuardian* works with *Squid*, making sure that all data passed to clients meets the criteria for what is acceptable.

Blocklisting

Guardian performs filtering of several different kinds, including URL and content not to mention MIME extensions, PICS ratings and others. Of course, it would be a bit much to



The web administration tool is rather elegantly put together.

expect the sysadmin to come up with lists of banned websites and regular expressions, so SmoothWall provides much of this information itself in the form of blocklists. These are updated about once a week and can be obtained over the web proving that the server is supported. One year's support is included within the initial purchase price, and then subsequent years must be purchased separately. The lists can be customised if required, and feedback such as recommending additional URLs to block can be provided to improve the quality of the blocklists in future.

One of the most impressive features of the product is the ability to do dynamic content filtering based on the text of a webpage. This is valuable because it neatly circumvents the problem inherent in URL-only filtering, which is that the domains used for various undesirable sites can be constantly changing. No image-scanning is performed though, so pages without text that can be analysed by the program could still slip through the net.

The major drawback with *Guardian* in its current form is that its user authentication is rather basic. There is no support for other directory services, meaning that the system cannot be integrated with, for instance, a Windows domain. The consequence of this is a maintenance overhead as a whole new set of usernames and passwords must now be administered. The alternative to this is to simply forget about authentication and rely on the connecting IP address to identify who has been going where. This might work in an environment of



Better to be denied access now than to have to explain why you were visiting dodgy sites later?

static addresses but is unreliable wherever DHCP is deployed.

All in all, *Corporate Guardian 3.0* is a solid offering, if not quite an essential purchase. Despite the name, it is perhaps unsuitable as a web content filter for big businesses, but is instead more likely to appeal to small- to medium-sized organisations. It works perfectly well and is reliable enough to be the backbone of any Internet access solution. Hopefully in future versions of *Guardian*, the *Samba* project's excellent *Winbind* program will be included and a major impediment to implementing this product will be removed. **LXF**

LINUX FORMAT VERDICT

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

An impressive product with lots of potential, but let down by limited support for authentication.

RATING **8/10**



Masters of Doom


Andy Hudson re-enters a world of nightmares...

BUYER INFO

- **AUTHOR** David Kushner
- **PUBLISHER** Piaktus
- **ISBN** 0-7499-2475-6
- **PRICE** £10.99
- **PAGES** 300

Where were you ten years ago? Chances are, wherever you were, you would have heard of or perhaps even played and completed *Doom*, the genre-defining first-person shooter. Unless you had been closely monitoring the bulletin boards of the day, *Doom* would have appeared to have come entirely out of nowhere for you. But in fact, the development of *Doom* and the fallout it caused was one of the defining moments in the lives of John Carmack and John Romero, two developers who – even today – are still at the forefront of games development. The release of *Doom* was such a huge leap forward in technology that currently, clones of it are still appearing, and every geek worth their salt knows what idkfa means (maximum ammo, armour and weapons – this cheat is most effective if you've found a backpack, of course). This book charts the story behind the technology, behind the people, behind the late night role-playing adventures and the almost addictive behaviour towards pizza and caffeine.

It starts fairly innocently with two rapid chapters giving us family backgrounds of the two Johns. Romero started his computing career at the local university, programming on a HP mainframe. For Carmack, it was a more humble beginning; he learnt how to use a TRS-80 at the local Radio Shack. The story picks up when the two Johns meet through work. They were both employed by a company that sent out subscription disks on a monthly basis. These disks were full of utilities and programmes (Perhaps a stone age coverdisc?!). Almost immediately they hit it off and started working together, convincing their management that they could do something similar, but with games. For those not brought up in the late eighties and early nineties, games at that point consisted of bland stick-like



sprites that moved in four directions. They believed that they could produce games like the world had never seen. Carmack was nothing short of a machine, even ending his sentences with a hum as if he was processing the information. Romero had a wild imagination that focused his energies in to making the best games he could think of. Underneath all this was a culture of adventure role-playing that coloured their lives.

Bright spark

The spark that set them on the course to producing *Doom* was ironically *Super Mario Bros*. This cutesy platform game was only available on the Nintendo and nothing like it was available or even possible on the (then) very basic PC. A conversation between Carmack and Tom Hall, a colleague, prompted an all night coding frenzy. At the end of it Carmack had produced a demo version of *Super Mario Bros* for the PC. Romero saw it and knew what it could mean to them. Releasing a game on *Gamers Edge*, the bi-monthly disk they produced, it quickly became popular throughout the bulletin boards in America. Scott Miller of Apogee wrote to Romero pleading with him to let him publish their games as shareware. Romero jumped at the chance, especially since things at the day job had turned sour. Moonlighting using the companies own computers (They borrowed them overnight), they

started work on what was to become a computing – not just gaming – legend.

Commander Keen was born through a lot of caffeine and pizza. It went on to earn the group called Ideas from the Deep thousands of dollars. Perhaps unsurprisingly, it didn't really faze them. They were totally focussed on their mission to create perfect games. Cheques for \$10,000 came and the only thing made of it was the fact that they could buy more pizza.

Castle Wolfenstein was a game that both Romero and Carmack had grown up with. When Carmack developed a way to create a first-person perspective in his games, both Romero and Carmack knew they had to use the *Wolfenstein* name. Luckily, the rights to the name were available and Ideas from the Deep snapped them up. What grew out of that was the now legendary *Wolfenstein 3D*, and a few months later *Doom* was born.

True, a lot of the imagery for *Doom* was clearly inspired by the all-time classic geek pastime of playing *Dungeons and Dragons*, but what the author of the book really drives at is the escapism that both men wanted to feel. Both felt their own lives were too boring and immersed themselves in the development and production of their games.

From a technical point of view,

there are facts and trivia scattered throughout the book. For the most part it details the specs of machines that were used in their careers, both of them starting their 'professional' work on the Apple II.

It also explains in layman's terms just how Carmack cracked the tricky problem of scrolling on the PC (Something that Nintendo had long done for their platform). From there, he talks about the texture mapping he developed for *Wolfenstein 3D* and how it would increase speed to have nothing on the floor or ceiling. His next development was the diminished lighting seen in *Doom* – giving us the creepy effect of walking through a corridor with a flickering light.

There are some anecdotes that raised a smile, such as when a certain university drop-out wanted to use *Doom* to promote Windows 95 at its launch – using the human rights-friendly tagline "Who do you want to execute today?" Other than humour and technical specs, the book is dedicated to its goal of helping you understand the thinking that drove two of the leading programmers in the nineties on to international acclaim – an entertaining read for both techy and non-techy alike.

LINUX FORMAT VERDICT

All you ever wanted to know about first-person shooters and the people behind them in all its gory detail.

RATING 9/10



Beyond Fear


Schneier has yet to put a foot wrong in his writing, and Paul Hudson was first in the queue for his latest.

BUYER INFO

- **AUTHOR** Bruce Schneier
- **PUBLISHER** Copernicus Books
- **ISBN** 0-387-02620-7
- **PRICE** £25
- **PAGES** 295

When we think of security, the first person to pop to mind is Bruce Schneier, long-standing security expert and field veteran. In the past, his books have all been best-sellers, and quite rightly – the mix of intimate personal knowledge of cryptography and other security methods mixed with his uniquely gripping style of writing makes for a blend few can resist.

In this, his latest book, he strays a little from his usual highly technical setting, which at first seems like a bad idea – the hackneyed “if it ain’t broke, don’t fix it” comes to mind. However, after reading just the first few paragraphs of the first chapter, you



realise you’re in for a rollercoaster of topics that are of the utmost importance in our lives today. Few authors are held in such respect that they can call the 9/11 attacks on the World Trade Center “amazing” and not be roundly criticised, and yet this is how the book starts – with the use of “amazing” in its truest sense, rather than the bastardised vernacular. From there, it moves onto the main point of the book: that all too often security is used to describe actions which often simply just limit our

personal freedoms.

The topics covered are largely computing-based, including “is online shopping safe?” and “will computerised voting machines make election results more accurate?”; and time and time again, the answer goes pretty much directly against popular belief. Naturally, Schneier backs all this up with logically structured arguments and copious references to external sources, and as usual, his arguments make perfect sense once explained thoroughly.

One of the biggest advantages of this book is that there’s not a single line of programming code in there – it’s simply well-written prose that picks up on day-to-day tasks and attempts to demystify them as part of a larger discussion. As such, you can pass this book to anyone, be they programmer, manager, or college student, and they’ll easily be able to get the most out of the book.

This is another winner from Schneier’s unfailing hand, and should be another addition to the bookshelf of everyone with an interest in security, computer-related or otherwise. If you’re really strapped for cash and can wait for all the insight given here, the paperback should be a little cheaper.

LINUX FORMAT VERDICT

Hard to put down and well worth the seemingly high price tag.

RATING **10/10**



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 “**FUTLIN049**”.

	KEY FEATURES	DISK SPACE	MONTHLY TRAFFIC (Gb)	SHELL ACCESS	PHP4.3.2 SCRIPTING	MYSQL3.2.3 DATA AREA	ADDITIONAL FEATURES	ACTIVE CGI LIBRARY	CRONJOBS	XML SUPPORT	PYTHON	FRONTPAGE EXTENSIONS	POP3/SMTP EMAIL ACCOUNTS	UNLIMITED ALIASES	EMAIL FORWARDING ALIASES	MEDIA FEATURES	HTTP STREAMING	REAL SERVER PROFESSIONAL	MONTHLY FEE*	YEARLY FEE*	SETUP FEE*
LINUX SILVER	100	1	Yes	Yes	0	Yes	Yes	Yes	Yes	Yes	Yes	5	Yes	Yes	Yes	No	£8.95	£89.95	FREE		
LINUX GOLD	200	4	Yes	Yes	1	Yes	Yes	Yes	Yes	Yes	Yes	10	Yes	Yes	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	£19.95	£179.95	FREE		
LINUX PLATINUM	500	10	Yes	Yes	2	Yes	Yes	Yes	Yes	Yes	Yes	50	Yes	Yes	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	£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

OpenGL Game Programming


Is this long-standing classic in the OpenGL world good for Linux users too, Paul Hudson wonders.

BUYER INFO

- **AUTHOR** Kevin Hawkins, Dave Astle
- **PUBLISHER** Prima Tech
- **ISBN** 0-7615-3330-3
- **PRICE** £43.99
- **PAGES** 777

When it comes to 3D graphics on Linux, OpenGL has always been the number one choice – it's well-supported by graphics card makers, easy to learn, and allows lightning-fast rendering of complex scenes. Yes, OpenGL does have its downsides (particularly extension management), but it makes up for this by being cross-platform and readily available for use in GCC.

As books go, once you're beyond the infamous Blue and Red books, one of the most enduringly popular books on OpenGL programming is this one – at 777 pages it's certainly not lightweight, and at £43.99 you're definitely paying a premium price for it.



However, a surprising amount of content is covered in there, including matrix mathematics, lighting and blending, multi-texturing, display lists, and particle systems. The content on particle systems is particularly well presented, and is pitched at such an easy level that everyone should be able to follow along.

For more advanced readers, there's quite a lot of content on difficult topics such as stencil buffers, quadrics, and shadow effects, although much of this is written from the perspective "This is difficult, so we're not going to go into

too much depth". On the upside, at least this gives novice readers a diving board into the topic, and it's impossible to cover *everything* in depth in this number of pages.

The key problem with this book is that it focuses on OpenGL as it relates to Windows. Yes, nearly all the graphical code samples can be used almost wholesale on Linux with little effort, but nowhere does it give the five or so lines of code required to make the examples *work* outside of Windows. For that you need an online

code site, such as NeHe's OpenGL tutorial (nehe.gamedev.net/) to fill in the gaps. This is a real irritation, because this is one of the best-written books for beginners on the topic, and the lack of cross-platform code is the only thing missing.

The book closes by creating a fairly simple 3D engine that allows you to load in a fully textured game world and interact with it, so it's very hands-on for newcomers. As a reference and general learning book for people wanting to get to grips with the essence of OpenGL and game physics, this is an excellent book. However, keep in mind that it will require a little of independent research to help you get the code examples working on your Linux system.

LINUX FORMAT VERDICT

Well written and covers everything for beginners, but too focused on Windows

RATING **7/10**



Text Processing in Python

Nick Veitch discovers snakes and strings do go together

BUYER INFO

- **AUTHOR** David Mertz
- **PUBLISHER** Addison Wesley
- **ISBN** 0-321-11254-7
- **PRICE** £44.99
- **PAGES** 514

Anyone who is familiar with Python will know it is a great and flexible language suited to a number of tasks. Text processing is certainly one of these, so it isn't as much of a surprise as it might be to see a book taking this specific approach.

This book doesn't attempt to teach you Python, but the thirty-odd pages of Appendix A do cover the basics of the language. In reality it isn't that complicated to understand, so this brief introduction may be enough if you only want to follow the examples laid out in the book.



The subject area is broken down into conceptual chunks, covering areas such as basic string operations, regular expressions, parsers and – most usefully – Internet uses such as email, HTML and XML. Much of the book is broken down into a series of problems, to which an answer in the form of Python code and an explanation is provided. This

makes it not just a good learning tool, but a great Python recipe book for problems you may come across. There is some particularly useful code discussed featured in the chapter on regular expressions.

Perhaps one of the most topical areas of discussion is XML. Although this is only covered in around ten

pages of the book, it does explain the different methods of manipulating XML data and provide some great examples which can be adapted for your own use.

The appendices cover not only the rudiments of Python, but also data compression, and Unicode, which is another forte of Python.

Overall, this book has the feel of a highly practical work. What you have here is not some grand work of theory, but plenty of real code that does real work in a real world. If you need to manipulate text on a regular basis, or want to find out how, buy this book! **LXF**

LINUX FORMAT VERDICT

Well written and with plenty of practical real-world application.

RATING **9/10**



Roundup

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

DVD Players

Tom Wilkinson takes time out from staring at a cinema screen to review six DVD players.



OUR SELECTION AT A GLANCE

- Xine
- Mplayer
- Ogle
- VLC Media Player
- Kplayer
- Xmovie

COPYRIGHT INFO

Film and TV show DVDs

- P35** *Sleepy Hollow* © Mandalay Pictures, LLC and Paramount Pictures
- P35** *The Man Who Wasn't There* © 2001 USA Films, LLC
- P36** *Angel* © 2002 Twentieth Century Fox Home Entertainment
- P36** *Lord Of The Rings: The Two Towers* (Preview footage) © 2002 New Line Home Entertainment, Inc
- P37** *An Association* © 2003 Jon Topper
- P37** *Red Dwarf III* © BBC Worldwide 2003
- P38** *Amelie* © 2001 UGC Images, Tapioca Films, France 3 Cinema

Not so long ago, it was impossible to find a DVD player application, or indeed any kind of easy-to-use media player to rival those found on other operating systems, such as Windows Media Player, or Apple's Quicktime Player. Happily, though, as Linux has become established on the desktop demand for such applications has grown, and today there are a variety of applications, all able to play DVDs straight after installation, with little or no extra configuration needed.

It hasn't always been quite as easy, though – until October of 1999, it was completely impossible to play most DVDs without purchasing a licensed DVD player application. These were of course only available for the proprietary operating systems. This came to an end with the publication of a small library called *DeCSS*, which enabled the user to bypass the Content Scrambling System (CSS) used to encrypt most commercial

DVDs. The development of this software was not, however, without its own risks. Development of software to subvert anti-piracy measures (including CSS) is illegal in the United States, among other countries, and Jon Johanssen, developer of the package, eventually stood trial in his native Norway accused of aiding piracy. He was eventually acquitted, though later others who had helped to distribute the software in the US were served with injunctions demanding the removal of the DVD producers' "trade secrets". Although the debate over what is and isn't a trade secret continues, most DVD playing software has now sidestepped the issue by hosting their web and download sites outside of the US.

While even a year ago, it was common for the applications to require a plugin before they would play encrypted DVDs, often downloadable from a site linked to from the players' own pages rather

than directly, this has changed. In fact in the majority of cases, the binary packaged versions of the software required some descendant of *DeCSS* to even install. *DeCSS* also appears to have the advantage that DVDs from any region can be played without having to resort to altering the drive's firmware, though this could just be a peculiarity of our test equipment.

On test are a variety of players designed to work in various environments – some will even display the DVD rendered into ASCII for those times when you need to watch a film but only have a video card capable of displaying text. Why you'd want to do this is beyond us, but there must be some demand for it.

Gimmicks aside, though, there are a number of features which a DVD player needs before it can be considered useful apart from CSS decryption. These include menu support, which makes it far easier to navigate through a disc to find the

relevant chapter, feature or extra; and full screen support allowing playback without the distraction of window borders and anything else that might pop up.

In addition, it's useful if your player does more than just play back DVD discs – some of the players on test here will play other types of video of assorted types, ranging from plain MPEG-encoded video to more proprietary multimedia formats such as Windows Media.

Lastly, the performance of the player needs to be considered. While the average desktop PC is more than adequate to play discs nowadays, some players can still show jerky video or muffled sound because they're still optimised for older hardware, and so are down-sampling (effectively throwing away) some of their input.

A player which combines the above features with a strong GUI will go a long way. Each of the players on test addresses some of these points.

Xine

Polished player with menu support.

■ **VERSION** 1.0RC2 ■ **WEB** <http://xinehq.de/>



Xine can handily skip through copyright notices and lead-ins – brilliant!

Supporting a variety of formats out of the box including SVCD, DVD and MPEG, *Xine* also has full support for DVD menus through the use of *libdvdnav*. This was initially developed

for use with *Ogle*, one of the other DVD players on test here, and demonstrates one of the best points about Open Source software – that multiple projects can share useful

code, benefiting them all. This isn't something that could happen with standard proprietary software!

The program comes in two parts – the core libraries, which provide the video decoding functionality. On top of that, it's necessary to install a front end. While we tested the package using the standard *Xine-ui* front end, there are a great many others available including GNOME and KDE applications as well as one for use with *Xine* in a set-top-box environment such as a digital video recorder. Three skins come with the player by default.

Playback of discs is extremely easy – if you have *libdvdcss* installed it's a matter of pressing the button marked 'DVD' on the user interface. If it's the first time that you've played the disc, *libdvdcss* will then start to try to crack the encryption on the disc, though this takes a couple of minutes at most.

Full-screen mode superimposes the controls on top of the movie display – but they're easily hidden with a mouse-click. One annoyance, though, was that the GNOME toolbars also stayed on screen while the *Xine* UI was present. This is most likely due to a limitation of X rather than GNOME or *Xine*, though.

The first time *Xine* is run, you're asked which of four levels of expertise you have so that the relevant options can be displayed. Even the beginners' menus are comprehensive. We had difficulty working out what even half of the options in the more advanced menus were for – the amount of tweaking the package allows is extreme!

Playback was smooth and the sound quality was crisp, even at hi-res, including deinterlacing of the picture, something which can cause problems with some players. Deinterlacing is necessary because DVDs are intended for view on television – a computer at a higher resolution and refresh rate won't always display the image properly.

LINUX FORMAT VERDICT

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

Very nearly the perfect DVD player for Linux, *Xine* is stable, easy to use and well-documented.

RATING **9/10**



Kplayer

KDE-Based media player based on *MPlayer*.

■ **VERSION** 0.31 ■ **WEB** <http://kplayersf.net/>



Kplayer is meant to be more user-friendly than MPlayer. Eventually...

KPlayer is still a relatively young project, reflected by its low 0.4 version number and as a result, there's still quite a bit of polishing required before it can

catch up with its more advanced cousins. *KPlayer* is an app that borrows heavily from the code in *MPlayer*. This is definitely not a bad thing though – a

solid code base means that authors can concentrate on other things. Sadly, there's some way to go with this. We had problems installing *KPlayer* from the supplied binary packages – they'd been built against a specific bleeding-edge version of the Qt libraries. Compiling the program from source, however, worked well enough.

However it shares *MPlayer's* limitations, too, like the lack of DVD menu support. Getting DVDs to play was also difficult – the docs suggested that the resource indicator `dvd://1` should be entered into a network dialogue box, which didn't seem to exist. While *KPlayer* will accept the location of the files to be played on the command line, along with the lack of menu support, this makes it difficult for new users to play discs.

One of *KPlayer's* biggest problems was that it seemed to crash consistently when right-clicking on the video playback panel. This is quite a major flaw, and one that I'd have expected to be fixed before now, even though the project is still less than a year old. As is always the case, though, this might just be peculiarity of our test machine.

The user interface is also minimal – no playback controls are present. At very

least, a video player needs stop, play and pause buttons, and preferably a slider to seek to given parts of the video. Those controls which do exist (by way of keyboard shortcuts) don't always work properly, the pause button in particular was problematic, and apparently paused the video while actually continuing to play it in the background (with sound still coming out of the speakers). The lack of preferences is also crippling – in the present version, the only preferences available are those of what key performs what function. As there's little other documentation, this also doubles as basic instructions. Given the lack of any features whatsoever, this screen is extremely important.

LINUX FORMAT VERDICT

FEATURES	3/10
EASE OF USE	2/10
DOCUMENTATION	2/10
PERFORMANCE	6/10

A player with a solid code base, but with many major issues to work through before it reaches stability.


RATING **4/10**



MPlayer

Capable player with support for many formats.

■ **VERSION** 1.0-Pre2 ■ **WEB** www.mplayerhq.hu/



MPlayer's default interface is command-line based, with a graphical frontend and compile-time option.

By default, **MPlayer** is a command-line app with no GUI. While this can be added at compile-time, it comes across as a late addition to the app rather than something integrated; there are skins available to download. *MPlayer* doesn't

support DVD menus: The developers suggest in the docs that if you want it, you should code it in yourself. The actual reasoning is more complex, and relates to the way the menu code in other apps often has to guess what it should be

doing. It does have the advantage, though, that unlike some of the other players on test here, it will play encrypted DVDs out of the box without requiring additional libraries. The developers decided early on to add additional functionality to their CSS implementation such as caching. This means each DVD needs its encryption cracking only once, making startup times for the second and subsequent plays much shorter. Without the menu system, though, this is only of limited use.

There are other advantages to counter these issues, though – *MPlayer* can be configured to use the binary codecs used by Windows apps to play media, and so can play things like Quicktime movies and Windows Media audio files – something that's only otherwise possible with Codeweavers' *Crossover Plugin*, or by running the respective players under WINE, which inevitably gives variable results.

In addition, *MPlayer* supports playback on a number of output devices in addition to the standard X11 output plugins. One of the ones I found most surprising was one that allowed graphical playback on a plain text console if you had an nVIDIA graphics card. While our test machine did indeed have such a thing (a GeForce 3), we were unable to

get this feature to work. As it's a new addition, it no doubt needs tweaking.

Playback on screen is smooth, with few jitters, and the sound quality is adequate. One thing that was surprising, however, were the frequent displays of a message stating the test system was too slow to play the video, and suggesting ways of improving the performance of the player. We found it curious that this was still occurring when playing back an uncompressed DV file on a 2.4GHz PC, but there was no other evidence of any problems.

There was also no choice as to which video on a DVD was to be played. While the first title is often the feature, this isn't always the case, and it makes selecting and playing episodes of a DVD of a TV show a bit troublesome.

LINUX FORMAT VERDICT

FEATURES	6/10
EASE OF USE	5/10
DOCUMENTATION	8/10
PERFORMANCE	8/10

A solid application providing support for a wide array of video and audio formats, but lacking a solid GUI and support for menus.

RATING **8/10**




VLC Media Player

Cross-platform player supports two-way streaming.

■ **VERSION** 0.6.2 ■ **WEB** www.videolan.com/vlc/

Formerly known as **VideoLAN Client**, *VLC* was renamed to reflect its functions beyond playing and rebroadcasting video streamed over a network. *VLC*

was simple enough to install – the download page also includes links to packages for each of the program's dependencies, which included *libdvdcss*,



An uncluttered app great for use on lower-end hardware, Mac OS X and *nix.

meaning that it supports encrypted DVDs without additional requirements. In fact, it was the first player on the market to support this. In addition, *VLC* uses the code originally written for *Ogle* for processing DVD menus.

The interface is clean with just play, stop, skip forward and backward buttons – built using the WxWindows cross-platform toolkit on top of GTK. Clicking on the 'Open file' or 'disc' buttons will allow you to play a disc inserted into the drive without mounting it, or to play any video file. There's also a button for opening a playlist, but the playlist editor is very simple and isn't that useful.

While not as comprehensive as *Xine's*, *VLC* still offers a wide variety of options to tweak the performance of the player in an easy-to-navigate tree.

Once a disc has been loaded, additional tools appear below the standard controls: a progress slider, allowing skipping to anywhere within the disc, and title and chapter skip buttons.

Full screen mode is supported, but uses the fallback of hotkeys – **Space** to pause, **F** to switch between full-screen and windowed mode, and so on.

One thing that was noticeable with the *VLC* playback was that the sound seemed to be less clear than that given in the output of the other players on test. It's unclear whether this is due to limitations in the player – it might be down sampling the sound input it's getting before outputting it for some reason, or using a different decoding method that's faster but sacrifices quality. Given the program's performance playing discs on slower machines, I'd guess at the latter. We were unable to find a suitable option in the preferences screen that allowed the audio quality to be increased.

LINUX FORMAT VERDICT

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

A solid player that offers a good deal, and will even work with MS Windows!

RATING **7/10**



Ogle

A dedicated DVD player application using GTK.

■ **VERSION** 0.91 ■ **WEB** www.dtek.chalmers.se/groups/dvd/



A good DVD player, but sadly lacks support for other media types.

Ogle was the first DVD player app for Linux to support DVD menus, through the development of the *libdvdnav* library.

Like *VLC*, *Ogle* has a simple GTK-based interface showing basic functions in its small control window. As it is a dedicated DVD-playing application,

though, it includes a few more functions including some for navigating through menus, as well as buttons for chapter skip and selection of subtitles and which audio stream should be used. Opening a DVD is easy enough, and the "Open Disc" menu option can be used, or just the "Play" button on the interface.

The menu option for the program's preferences don't appear to do anything, which begs the question as to why it was included. There's also no GUI while playing back in full screen mode, which makes it difficult to control. Switching between modes can be quite annoying.

Ogle has a number of drawbacks compared to players such as *MPlayer* and *Xine* – it supports only one media type. This approach does have some advantages, though. It means that the authors can aim to write the best DVD player, ignoring other formats. The relatively slow update cycle of *Ogle* compared to that of bigger projects means that it's not necessarily able to keep up with the development pace of its rivals – looking at the program's web page implies that updates happen around once every three months or sometimes even less.

Performance, however, is good on low-end hardware, giving a jitter-free picture and crisper sound than *VLC*. There's no support for deinterlacing, though, which results in some problems at high resolutions, with video appearing jerky. This can be fixed to some extent by reducing the resolution that X is running, but it's still a minor annoyance.

Ogle also comes with a command-line interface, which optionally supports text-based rendering using *adlib*.

Gimmicks are nice, but it would be more beneficial to the project to work on more useful features, such as deinterlacing the output picture, or full-screen mode GUI.

There are few options that can be passed to the *Ogle* binary – just whether to start in text or graphical mode and which file or disc to open. The manual pages do include large amounts of documentation on how to configure the player using its XML-based configuration file. It's still not nearly as easy to configure as it should be, though.

While *Ogle* has all these drawbacks, it's still good at what it purports to do – play DVDs. It could be argued that the lack of options and formats supported makes it very easy to perform the limited functions that it does work with, so makes sense for new users. For the rest of us, though, it would probably better to use *MPlayer* or *VLC*.

LINUX FORMAT VERDICT

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

A highly limited player, but one that's certainly good at what it does.

RATING **7/10**




Xmovie

Lightweight app intended for uncompressed video.

■ **VERSION** 1.9.10 ■ **WEB** <http://heroinewarrior.com/xmovie.php3>

Originally written as a replacement for *Xanim*, this player very much shows its age, unfortunately – it's built with the

Motif toolkit, which was common before GTK and Qt, used to build GNOME and KDE respectively. This means that it



It's potentially a great viewer for HDTV but not DVDs in its current state.

doesn't fit in to any modern desktop particularly well, with a dark green menu bar and light green video seeking slider.

The player is, however, well designed. The controls at the bottom of the screen give what is required – at least that's the theory. In practice, selecting any of the controls during playback appeared to cause the entire app to freeze, requiring it to be killed from the command-line. Switching between full screen and windowed mode worked well, though, and didn't cause problems.

There are also options for sound mixing – the player can down mix Dolby Surround into stereo, or even up mix Stereo into Dolby, though it's hard to say what use that would be. Prescaling and cropping of the image was also supported. Any number of options could be changed while the video was playing, without any problems at all, surprising given the problems with basic playback functions. Options are forgotten between playbacks, though, which is irritating.

The DVD playback aspect of the program is unfortunately lacking, too – only unencrypted DVD playback is

supported, and there's no support for either menus or chapter selection. To play a given DVD title, it needs to be mounted and the individual media files need to be opened in the player. This makes it even more difficult to work out which file is the one required than it is when using *MPlayer* and *KPlayer*, which rely on the user entering the correct track number in order to play the video.

The docs comprise only of a list of formats supported by the program and a short description of the menu items, most of which are self-explanatory. More info, such as how to open a DVD movie (it's far from obvious) would be nice.

LINUX FORMAT VERDICT

FEATURES	4/10
EASE OF USE	5/10
DOCUMENTATION	3/10
PERFORMANCE	6/10

A simple-to-use media player but not really one to consider when playing DVDs

RATING **3/10**



DVD PLAYERS THE VERDICT

As time goes by, the selection of DVD players for Linux emerges as increasingly able to topple the commercial players available for other OSes. Long gone are the days when a reboot into Windows was required to watch a disc on your PC. Also, happily, long gone are the days when it was difficult to play region-coded discs, many of the players officially supporting playback of such things despite legal ambiguities. This would have been unthinkable even two years ago.

Almost all the players on test performed admirably. Of the two low points, *Kplayer* did well enough but is still early in its development cycle – no doubt it will improve in leaps and bounds over the next few months. Unfortunately, the same can't really be said of *Xmovie*, which at present doesn't seem sure what it wants to be – it's not really suitable for the playback of DVDs, and seems better geared toward broadcast media, something which *VLC* is designed to cope with from the ground up.

Code in common

A good deal of code sharing has gone on between the applications as well – *libdvdcss* was initially developed for use with *MPlayer*, but is now used by almost all the players in one form or another. The menu navigation code originated with *Ogle* but has since been integrated into *VLC* and *Xine* as well, as have the chapter location codes. This spirit of co-operation has no doubt improved the quality of everyone's apps and improves the choice available to the end user.

Performance-wise, there's not a lot to choose between the differing players – *VLC* had slightly worse sound quality,



***Xine's* total support for the entire DVD disc, including subtitles and menus means that it comes out on top.**

but not so much as to affect enjoyment of the disc being played back. *Ogle* showed slightly worse picture quality, but this was probably more to do with the way in which it handles interlaced pictures than anything else.

Both *Xine* and *MPlayer* come with interfaces suitable for use in a set-top box environment, and either would be a worthy addition to the distributions that are aimed at that very market. In this corner, though, *Xine* has to come out on top due to its support for menus, making finding the video clip you wanted to play back very easy indeed.

MPlayer does have the edge in other arenas, though – more output options than any of the other players, and thanks

to the way it can handle Windows plugin and playback libraries, plays back far more types of video, including those that are exclusively licensed to certain applications, such as Apple's Sorenson Quicktime codec, and Windows Media. The main point against *MPlayer*, and the reason it hasn't scored higher, are the lack of menu support, even two years after other players have had it, and the lack of a default graphical interface. If these two points are ever fully addressed, *MPlayer* will no doubt overtake the other players quite significantly – it's already overtaken the Linux kernel as the most popular project available from the open source software repository Freshmeat (*Xine* comes 5th in the same chart).

VLC and *Ogle* are both fine players too – this is far from having been a two-horse race, though each of them is lacking some of the features present in the other players, and you could do worse than use either. In the end, though, only one player could cope with everything that we'd want from a DVD player, and that was *Xine*. Having borrowed the best features from each of the other projects, allows it to combine them all in a package that gives easy access to almost everything required from a media player. Those gaps are filled admirably by *MPlayer*, so we'd recommend having both installed on your PC, neatly covering all the potential bases. [LXF](#)

TABLE OF FEATURES

NAME	FULL SCREEN	CSS SUPPORT	MENU SUPPORT	GUI TYPE	SKINNABLE	COMMAND LINE SUPPORT
<i>Xine</i>	Yes	Yes	Yes	Standalone	Yes	Yes
<i>Kplayer</i>	Yes	Yes	No	Qt	No	No
<i>MPlayer</i>	Yes	Yes	No	Standalone	Yes	Yes
<i>VLC Media Player</i>	Yes	Yes	Yes	GTK	No	No
<i>Ogle</i>	Yes	Yes	Yes	GTK	No	Yes
<i>Xmovie</i>	Yes	No	No	Motif	No	No

HotPicks

The best new open source software on the planet!



Mike Saunders

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

This is the place where we get to profile some of the hottest software around.

Each month we trawl through the hundreds of open source projects which are released or updated, and select the newest, most inventive and best for your perusal. Most of the Hot Picks are available on our coverdiscs, but we've provided web links if you want to make sure you have the very latest version.

If you have any suggestions for things that we should cover, email us at linuxformat@futurenet.co.uk

HOT PICKS AT A GLANCE

Kile	40
CGI:IRC	41
Meld	41
AntiRight	42
Gmodconfig	42
APT-Fu	43
NewsBruiser	43
Enigma	44
GL-117	44

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!



KDE LATEX ENVIRONMENT

Kile

■ VERSION 1.6 ■ WEB <http://kile.sourceforge.net>



Kile's main editing window in action with the Quick Start wizard up front.

TeX, the 'tau epsilon chi' document typesetting system – commonly pronounced as 'tech' – has a lengthy history behind it. Professor Donald Knuth created *TeX* to assist in preparation of his book *The Art of Computer Programming*, and made it freely available for all to use – today, it enjoys popularity among the scientific, technical and academic spheres thanks to its versatility in producing highly complex documents.

TeX (and add-on *LaTeX*) files are built up from commands and text (in a roughly similar way to HTML) and the emphasis is on structure and layout rather than WYSIWYG editing; however, the friendlier *LyX* program provides a more hands-on approach if required. For more info on the system and its uses, dig out *What On Earth is TeX?* article from way back in LXF01.

Kile's goal is to be a an approachable *TeX* and *GnuPlot* front-end built around the KDE desktop suite, with plenty of features to keep experienced users happy while making the editing job more straightforward and

speedy. You'll need KDE 3.1 to build this release from source, along with numerous *TeX* bundles and the *xfig*, *transfig* and *gnuplot* packages as well. We've supplied RPMs for Red Hat 9 and SUSE 8.2 on the coverdisc – these may work on other distros, as long as the other requirements are installed.

With a slightly cramped layout in lower resolutions, *Kile*'s main window holds a bunch of rearrangeable toolbars along with the navigation section and main syntax highlighted editing component. Beneath this sits the tabbed output box which includes an embedded *Konsole* for quick command line work. It's an appropriate setup with good icons and plenty of tweakability, and tooltips abound to make the first few uses generally hassle-free.

Joy of TeX

When creating a new document, *Kile* offers a handful of ready-made templates (plain, article, book, letter and report) and new templates can be defined from other files. Multiple documents can be opened and

worked on at the same time thanks to the tabbed editing box, and quick access to the most commonly used formatting commands (bold, italic, alignment and font size) is available via the toolbar. Most menu operations have redefinable keybindings, including wisely chosen **Alt+N** keys for quick building and viewing, and the essential cut/paste, find/replace and multiple undo are all present too.

Kile sports various viewers built in to the main program – DVI, PS and PDF formats all catered for – and external viewing tools can be chosen if need be. (By default, the program builds with *LaTeX*, then converts to PS with *dvips* and runs a viewer.) Along with menu access to many frequently used *TeX* commands, new user-defined commands can be added inside the program (automatically given a **Shift+Fn** keybinding) together with additional tools. These features in particular will be welcomed by more demanding users.

A few screenshots improve *Kile*'s skimpy but pretty much adequate online documentation, although some more information on the workings of the *GnuPlot* front-end would be well worth throwing in. Similarly, the program options are a little thin; autosave duration, file conversation commands and editor/spelling checker settings can be tweaked, among others, but the excellent keybinding and user-menu configurability more than makes up for this.

All things considered, *Kile* does a first-rate job as a slick and polished *LaTeX* editing suite. The impressively well constructed interface does eat into screen real estate at anything lower than 1024x768; still, being able to move its components around and configure the toolbars stops this being too much of a problem. Clearly *LyX* is a more suitable choice if you're looking for easy-going near-WYSIWYG editing, but, armed with a *TeX* tutorial and the online documentation, most newcomers will be comfortable with *Kile*. Its featureset for long time users makes it a superb all-rounder too, and we recommend it to all regular *TeX*ers.

WEB TO IRC GATEWAY

CGI:IRC

■ VERSION 0.5.3 ■ WEB <http://cgirc.sourceforge.net>

Internet Relay Chat is a mature and proven system for real time online discussion, combining user application choice with a well-defined and open protocol. In contrast, Web-based chat rooms tend to be complicated affairs to set up and run. Nevertheless, it's difficult for site admins to expect their users to learn about IRC servers, channels and /FOO commands, so thankfully LXF reader (and Perl guru) David Leadbeater's CGI:IRC is a fine solution. If you're impressed with this program, contact him mentioning that you got his app from *Linux Format* – user feedback is always welcome.


CGI:IRC helps web server admins provide a gateway to IRC through the browser. In this way, there's no need to ask your visitors to download a client and read the docs – you just point them at the relevant page and they're

off. Setting up is a breeze; with the configuration file edited (selects default server, theme and sets restrictions), the files need to be uploaded to the cgi-bin/ dir and that's it.

Apache 1.3x or *2.0.44+* and *Perl 5* are required. Supported browsers include the *Gecko* gang (*Mozilla*, *Firebird*), *IE5* and *IE6*, *Konqueror 3* and even *Links*, but a few minor features may be missing with some. A user points his/her browser at `cgi-bin/irc.cgi`, and is presented with a login box – one more click and a familiar IRC client layout appears, crafted together with frames and driven by JavaScript.

Rabbit, rabbit

Buttons in the top pane control channel selection, while private chats can be initiated with a double-click on the user list. CGI:IRC's stock glass icon



Chatters nattering away in a #debian @ freenode session (with login box displayed on the right of the screenshot).

theme (changeable) is clean and easy on the eyes, and in general it operates just like any basic IRC client with a praiseworthy amount of functionality considering the limits.

At this point, the only major issue is that each user running CGI:IRC needs an individual Perl process and

consequently munches up around 4MB of RAM. The author hopes to resolve this in a future version, but for now it's a clean, well-designed and very competently coded choice for those who need a no-nonsense chat room solution or are behind an ultra restrictive firewall.

VISUAL FILE DIFFERENCES APP


Meld

■ VERSION 0.9.0 ■ WEB <http://meld.sourceforge.net>

Without doubt, the 'diff' utility is one of the most helpful and commonly used tools in existence for UNIX programmers – it reads through similar files and plucks out the differences, allowing any number of coders on a project to spot all the changes. Additionally, the resulting output can be fed back to the original file(s) as a patch, so end users can update the source code from, for example, kernel 2.4.21 to 2.4.22 with ease. Change descriptions are much smaller than new source downloads too.

Although *diff* is a command-line tool, sometimes it's more useful to view changes in a graphical environment; *Meld* does this via a Python/GNOME 2.0 interface, offering a colourful indication of what's new or different. Along with

the GNOME 2 libraries and headers, you'll also need the Python bindings



Meld displaying a few changes to a C source file – note the coloured margin boxes highlighting what's new.

for *Glade*, *GConf* and *Canvas* (typically called *pygtk2.0-libglade* and *gnome-python-gconf* etc in distros that are RPM-based).

Meld's main screen includes a multi-pane setup for showing the file differences along with a toolbar for common actions. To start off, you're prompted for filenames (two-way or three-way *diff*) and it's always great to

see a tabbed interface for working on multiple files simultaneously. A blue highlight shows modified bits, while red on green is used for new/missing sections, and the stretchable lines between panes give a quick indication of how it all fits together. Smartly implemented, then, and the ability to edit files on the fly with unlimited undo/redo and spot the changes immediately is a great help.

Directories

Along with text files, *Meld* can pick out file changes between directories (particularly handy if you're comparing with an old backup) and includes CVS functionality for finding updates and committing changes. There's little in the way of options and the online documentation is brief, but that's not a problem as it's all straightforward to use. Planned for the next version (built around GNOME 2.2) is syntax highlighting courtesy of *gtksourceview*, and it's already a reliable and helpful little application for coders.

LOW RESOURCE DESKTOP SUITE

AntiRight

■ VERSION 1.97 ■ WEB www.nongnu.org/antiright/


Before **GTK+** and **Qt** established themselves as the main graphical toolkit choices on Linux systems, *Motif* was the *de facto* standard with commercial apps such as *Netscape*, *WordPerfect* and many others built around it. *LessTif* provided a completely free implementation of the toolkit early on; now with *Open Motif* being included in most recent distros, we're starting to see a resurgence of projects using the widget set.

Reader Bob Holden suggested we look at *AntiRight*, a small desktop environment with a focus on easy operation for newcomers, low system requirements and a sprinkling of features for power users. You'll need either *LessTif* or *Open Motif* to build from source – it compiled without any hoop-jumping on *Open Motif 2.2*, and the only pre-build configuration required is a quick edit of *config.template*.

AntiRight doesn't include a window manager of its own, but the perfectly

appropriate *mwm* (*Motif Window Manager*) is included with the requirements mentioned above. Firing up the *antiright* binary with the WM started opens the main toolbox, which provides quick access to the suite of small utilities – no icons, no fancy frills. Its inbuilt calculator and command line are useful extras. Among the supplied tools are an unsophisticated text editor (standard facilities but no line number indication *etc*), a very basic file manager, an analogue clock, a process list and a disk mounter. Other buttons can be user defined to external mail/browser apps, and on the whole it's a usable setup without any trimmings. With few *Motif* applications in common desktop use today, it can look out of place, but *GTK+* and *Qt* both offer *Motif*-like widget themes for consistency.

Clearly, *AntiRight* doesn't compare to the major desktop environments of KDE and GNOME but it's enormously less demanding on memory and CPU,



AntiRight's bundle of tiny applications with the no-nonsense *mwm* running as the window manager.

and therefore ideal for an old 486 or Pentium 1 box. Coupled with *Dillo*, *Syphed* and other low-resource apps it's capable of day-to-day tasks;


equally, the easy going little toolset and uncomplicated operation make it ideal for a child or Linux novice taking his/her first steps in computing.

KERNEL MODULE CONFIGURATOR

Gmodconfig

■ VERSION 0.5 ■ WEB <http://gmodconfig.sourceforge.net>

Back in the early days of the Linux kernel, loadable module support didn't exist and every required driver and feature had to be compiled right in to



The *ov511* module is the only one fully documented at present.

the kernel image. This wasn't a problem for experienced admins, but it made life difficult for first-timers; equally, distro engineers had a hard time working out what users would need out-of-the-box. The module code today works well, but still requires some manual tweaking in places – *Gmodconfig* aims to make all this a great deal easier through a graphical front-end.

As it's built around GNOME 2, you'll need the development libraries if you're building from source. We've supplied an RPM package on the coverdisc which should work with most recent distros (you may need to make symlinks to *librpm-4.1.so* and use **--nodeps** if you're using a more recent RPM version).

Options

On the main screen, *Gmodconfig* sets up a tree list of modules for the currently running kernel, while clicking one brings up an options pane on the right. This takes the information returned from *modinfo* and presents it in a more graphical form – that's

decent enough, but the real strength lies in the XML description files. Providing far more information than *modinfo's* output, these XML files can offer in-depth option text in multiple languages along with proper checkboxes and drop-down lists instead of blank fields.

This makes generic driver configuration orders of magnitude easier for newcomers (distro-specific tools aside), and there's early code for an online updating system in place. Unfortunately, there's only one description file included at present but a tool has been thrown in to create new XML files, so hopefully the range will grow over time.

If an active surrounding community is built up, *Gmodconfig* could make a revolutionary change to the way end users manage and configure their kernels, and module repositories for the major distros would make new hardware support just a click away. There's still plenty of work to be done but *Gmodconfig* is an intriguing project with enormous potential, and one to watch closely in the future.

DEBIAN SOURCE PACKAGE BUILDER

APT-Fu

■ VERSION 0.2.2 ■ WEB www.yhbt.net/normalperson/debian/

Building software from source rather than relying on pre-compiled binary packages is all the rage at the moment – mostly thanks to Gentoo, which has demonstrated how a source-based packaging system can work without too many troubles. While the differences aren't always huge, it's certainly true that compiling for a specific architecture can lead to performance gains in many apps and provides more scope for tweaking build options. *APT-Fu* is a small program that intends to make source compiles on Debian more elegant.

Installing *APT-Fu* is a simple case of `dpkg -i` on the package, and its few dependencies can be grabbed easily from the current Stable archive. During the first proper run, the program will pull in the essential .deb building assistant packages (*debhelper*, *fakeroot* etc.)

automagically, so with a suitable `sources.list` it's ready to go from the off.

For basic usage, entering `apt-fu src-install <package>` will grab the package's source tarball and .dsc from the Net (or CD-ROM if specified in `sources.list`) and compile and install via the usual Debian process. Unless explicitly requested, *APT-Fu* will pull only binary packages of runtime dependencies although those too can be built from scratch if needed. It all works marvellously well and the program interjects the build output with its own messages as a reminder of what's going on.

With the help of *optFiles*, more detailed control over the build process can be applied – patches to the build scripts can be added, environment variables set and other adjustments made. A few are supplied with the package for *glibc*, *Firebird* and others, and more will appear in future.



APT-Fu's output from a 'src-install' of W3M, grabbing build dependencies. This application will help make Debian more accessible for the masses.

It's still a little rough around the edges, with the occasional error message popping up, but in our testing we encountered no major problems and the program works beautifully. We'd like to see *APT-Fu* (and the other


Debian source building tools) in the next Stable release; they'll do a fine job of joining Debian's renowned rock-solid stability with the control and performance tuning optimisations of Gentoo.

BLOG SERVER

NewsBruiser

■ VERSION 2.1.0 ■ WEB <http://newsbruiser.tigris.org>

As the whole Web-logging (or 'Blogging' as it has popularly become known in the last couple of years) phenomenon continues unabated,



Blogging's never been more pleasant – NewsBruiser's config panel on top.

stacks of tools and applications are in development to help users create fanciful diaries on the Internet. For those unfamiliar with blogs, they're essentially online journals in which a person posts opinions on topical subjects or feelings, and others can pop in and add comments. It's gigantically popular right now, and *NewsBruiser* is one of the best blogging utilities doing the rounds.

Mommy dearest

Striving to be easy to manage (the programmer's mother uses it!), *NewsBruiser* requires Python (1.5 or 2.x) and a working *Apache* server with SSI enabled. Installation is brilliantly simple – a CGI script copies the necessary files into place and performs some tweaking, and then it's all set to go for the first entry. The only manual intervention needed is some permission `chmoding`, and it's all detailed in the well-written docs.

In its default setup, *NewsBruiser* assembles a simple table based layout for the front page with a calendar on the right and entries on the left. A

search box is available, and an assortment of CSS 'themes' have been included for good measure; these range from merry pastel shades to dark and moody schemes.

When editing an entry, images can be uploaded onto the server and then included in the text with `%filename%`. Being able to mix in HTML is a bonus, although having a few buttons to insert tags (or indeed a quick tag reference) would improve things. Impressively, user comments can be monitored for spam (a growing problem) and the whole system runs like clockwork.

NewsBruiser offers hugely impressive flexibility in its configuration panels; templates, RSS feeds and trackbacks can be set up, while all manner of cosmetic and under-the-hood aspects can be fine tuned too. The generated HTML isn't too ambitious and works fine in all major browsers, and the default themes give an attractive and uncluttered appearance. Friendly, organised, versatile and great to use, *NewsBruiser* is well worth a look if you're running (or plan to run) a blog.

OXYD-INSPIRED PUZZLE GAME

Enigma

■ **VERSION** 0.81 ■ **WEB** <http://www.nongnu.org/enigma/>

Although blasting the brains out of hideous aliens is a lovely stress reliever, it rarely gets the grey matter moving and sometimes puzzle games are a good way to wind down (or get wound up). *Enigma* describes itself as a puzzler "inspired by *Oxyd* on the Atari ST and *Rock'n'Roll* on the Amiga," and is one of the best we've covered in *Hot Picks*.

When compiling from source, you will need to have recent versions of the *SDL*, *SDL_mixer* and *SDL_image* multimedia libraries (and their associated development packages) installed. You can find the most up-to-date versions of these on the coverdisc in the Essentials/SDL/ directory. With those in place, compilation and installation should go hitch-free.

Enigma's lavish presentation, a combo of smooth images and fitting title soundtrack, has evidently seen

much attention; as we've found before in *Hot Picks*, open source games can easily match commercial titles in terms of polish with a little effort. The crisply drawn sprites and attractive backgrounds don't get in the way of gameplay – a mistake made all too often in games of all genres.

The main objective is simple: direct a ball into blocks, revealing colours on the way, and then match those blocks. Well, simple in concept but screamingly difficult on later levels, where all manner of hateables including lasers and skull stones hinder your progress like a massive blancmange in a busy rush-hour road. Using the mouse to control the ball's motion really works splendidly (and, oddly enough, it works especially well on laptop TrackPoint devices too) with convincing momentum and a hypnotic click on contact.

Items can be picked up along the way, including springs and explosives,



Hell's bells – another fine mess I've got mesel' into. Yet again.


and the clever designs over 141 levels lead to some fantastically taxing puzzles. It's a shame the full-screen and window video modes are limited to low resolutions, but this can't be helped without ugly sprite scaling. All

in all, *Enigma* delivers a top-notch desktop distraction with an almost perfect mixture of fast reactions, hand-eye co-ordination and careful planning. It could prove to be hugely addictive, though.

ACTION FLIGHT SIM

GL-117

■ **VERSION** 1.0.1-2 ■ **WEB** <http://home.t-online.de/home/Primetype/gl-117/gl-117.html>



An aerial picture from the Lake District Tourist Board's new 'shock and awe' advertising campaign.

Thanks to *Top Gun* and other action-centric movies on that theme, the life of a fighter jet pilot is endlessly romanticised, and yet, curiously, few seem to consider that it's a ridiculously dangerous and complicated career to follow. But no, love affairs and rolling jeeps and resignations and final showdowns are the staple diet of such flicks, and as a result every little boy dreams of flying a terrifyingly fast aircraft through mountain gorges and exchanging ammo with the enemy. And thanks to *GL-117*, you can do it all at home now.

To build *GL-117* from source, you'll need a few dependencies which should be installed by default on most recent distros. Along with OpenGL/Mesa, *GLU* and *GLUT*, you should also have *SDL* and *SDL_mixer* and their relevant development packages in place. From there, compilation shouldn't bring up any nasty errors.


GL-117's thumping title soundtrack and soaring aircraft views open the game, with the main screen a hub for the various gameplay modes.

Choosing a Free Flight is the quickest

way to become accustomed to one of the three jets, and lets you roam the landscape without distractions – the most crucial controls are a doddle to master with the mouse and/or keyboard. *GL-117's* PDF manual gives a quick overview of aerodynamics, forces and axes too for the 'serious' player.

Aside from the flying and dogfighting tutorials, the game's Campaign mode is where most of the action takes place; here you're given a set task to accomplish, using your piloting skills and weapon mastery to destroy ground vehicles and battle enemy fighters in the air. There's a good mixture of missions and the gorgeous graphics, while not highly detailed, work well.

Ignoring the occasional spelling mistake in the mission text and the seemingly Caesium body of the plane (explodes on contact with water), as well as the few other refinements needed, *GL-117* is an entertaining action fest and handy for proving to nay-sayers that Linux can do 3D games just as capably as other OSes that don't bear a mention. **LXF**



on the desktop

Although it was late to join the Linux party, Sun Microsystems is poised to make a splash by taking Java to the desktop. **Paul Hudson** has a look to see what's brewing...

Since it was launched on May 23rd 1995, the Java programming language has gone from strength to strength, and is now recognised as the leading cross-platform development tool available. One of the leading reasons for its incredible success has been the carefully guiding hand of its creator, Sun Microsystems, when designing and implementing language changes.

Although two of Sun's slogans are "The network is the computer", and "We make the net work", Sun wasn't first off the starting blocks when it came to utilising the power of Linux as a network operating system. Having said that, it has never been a company that's slow to admit its mistakes and re-assess its position; and the first whispers of Sun developing an all-new desktop system were heard in Spring

2003. At Sun Tech Days 2003, Scott McNealy (CEO, Sun) made the first serious moves towards collecting support for the release of a new desktop system as he told 2,000 developers "We want you to build the next-generation software alternative to the Microsoft architecture".

From that date onwards dozens of rumours, most unsubstantiated, floated around the Internet about *Project Mad Hatter*, a new, if somewhat peculiarly named, initiative by Sun to develop this next-generation software. These rumours culminated on August 5th,

when, at the LinuxWorld conference in San Francisco, Sun previewed the Mad Hatter project for the first time and revealed many of the components that made it up. Powered by GNOME, *StarOffice*, *Evolution*, and *GAIM*, the *Mad Hatter* demo wooed the crowds and, with Java at its core, also started the ball rolling on what was soon to be known as Java Desktop System (JDS).

Caffeine high

With Java installed on over 550 million desktops around the world, "pervasive" is a word that can be applied to Sun.

"Having a predictable outgoing that's simply a calculation based on number of employees will save time and money."



As such, it's not surprising that Java sits at the core of Sun's new desktop push – or so you might assume. Based on Linux, an OS written in the C programming language, and utilising XFree86, GNOME, *StarOffice*, *Evolution*, and *GAIM* all written in C and C++, you may be forgiven for wondering exactly what part of the desktop system *does* run on Java.

The answer is “very little”, as the name has been chosen to leverage the Java brand, which many would say is much better-known than the Sun brand itself. Instead of forming the base for the rest of the system, Java is instead used for third-party apps, and this position makes good business sense – the simplicity of writing efficient, cross-platform apps in Java makes it a no-brainer for developers wanting to target the largest audience.

Once you understand that JDS isn't just about pushing Java for the sake of it, it becomes clear that Sun is actually in a remarkably good position to market a working, alternative enterprise desktop solution. Keeping in mind that Sun was co-founded by Bill Joy, the principal designer of BSD, that Sun already has in place a multi-tier technical support system for the thousands of companies that already rely on it, and also that Sun has many years of ‘big iron’ enterprise-level experience through its successful and long-standing Solaris OS, it certainly puts Sun in an enviable position to take full advantage of Linux.

Furthermore, Sun is using JDS to continue its vision of simplifying networked computing, and this is most

JDS AT SUNNETWORK 2003

Paul Hudson was there... News report next issue!

As nearly everyone predicted, the Java Desktop System launched with much ceremonial fanfare at the SunNetwork Conference in Berlin in December 2003, with a keynote speech by CEO Scott McNealy highlighting the differences between JDS and its competition. Referring to the program menu button, McNealy said “Notice that it's ‘Launch’, not ‘Restart’”, amongst other seemingly anti-Microsoft and anti-Intel comments.

Also highlighted was the first public airing of *Looking Glass*, a test user interface modelled entirely in Java 3D that has 3D windows, transparency, and smooth movement animation that Mac OS X only wishes it could achieve.

On the business front, Sun announced that the Java Enterprise System was

actually going to be available for free for companies that employ under 100 people. Furthermore, along with the other 6,000-odd delegates attending the conference, we got to see the first proper preview of the new Java Studio Enterprise, the successor to the current Java IDE tools made by Sun, that will cost just \$5 per employee per year if you also license JES.

Many of the expo attendees that we spoke to were more than happy to discuss their plans for porting their software to Linux, in line with Sun's policy of increasing its activity on the Linux front. We expect to hear lots of news about exciting Linux developments from Sun and their partners over the next few months.

clearly visible in its licensing policy. Traditional commercial licensing methods for software vary wildly, from a fixed cost-per-CPU, to client access licences (CALs), to hot-seat licensing – all difficult to track and manage – which is why Sun is offering a new way in JDS.

For a fee of just \$100 per year, a company can license JDS for one employee in their company, irrespective of how many computers that person has or where they are located. While this makes the life of

system administrators a great deal easier, it's also set to drastically ease the job of accountants the world over who have to forecast software costs – having a predictable outgoing each year that's simply a calculation based upon the number of employees will save a lot of time as well as money.

Doesn't cost a bean

Given that the retail cost of *StarOffice 7* is \$80 and the rest of the bundled software is available for free in one






JDS is positively oozing eye candy, an area that Sun clearly thinks that customers are looking for, and the login screen is perhaps the most strikingly attractive part of the system.



JAVA DESKTOP



The Launch menu is shorter than most of us are used to, which is no bad thing. You'll find the actual Java apps skulking away under the Extras submenu, though.



form or another, you may be forgiven for wondering how the \$100 a year price tag is validated. Firstly, the support on offer is impeccable – all subscribers get free maintenance and updates for their machines, as well as a full 60 days of installation and configuration support. Furthermore, Sun has actually done a lot of work to the system to hide much of the

complexity inherent in Linux, and the end result is quite impressive – more on that later.

This cost is lowered further if Java Enterprise System is also deployed. JES takes the same principles of JDS and applies them to the server side – for \$100 per employee per year, companies get Sun's directory server, identity server, proxy server, application

server, message queue server, web server, messaging server, calendar server, and portal server. That's a massive amount of software for a very low price, but even better is the option that if you license both JES and JDS, you get the lower price of \$150 per employee per year.

For companies looking to jump ship from a Microsoft solution, this pricing must be very attractive – one copy of Windows XP alone (at basic prices) costs US\$300, a copy of *Microsoft Office Standard Edition* costs US\$400, *Windows Server 2003* costs US\$1000, client access licences to connect Windows XP to *Windows Server 2003* cost US\$200 for five licences, *Exchange 2003 Standard Edition* costs US\$700, and client access licences to connect Windows XP to *Exchange 2003* cost US\$67 a licence. Excluding the purchase costs of *Windows Server 2003* and *Exchange 2003*, the total cost is around US\$800 per user – equivalent to over five years of licensing for both JES and JDS combined.

If you're thinking, "Yes, but after those five years, the Microsoft company will be better off, as the JDS/JES company will still have to continue paying their licensing costs", then you must consider that Microsoft no longer offers any support for Windows NT, its operating system of five years ago, which means that any

DUKES OF HAZARD

Time to ask the big question: how much is Sun betting on JDS?

While companies like IBM and HP leapt onto the Linux bandwagon, Sun was a great deal slower in making its mind up. Depending on who you spoke to, or even the day you spoke to them, Sun's view of Linux would vary from full support to disparaging comments.

For example, Scott McNealy said on September 16th "We have got SuSE and Red Hat on the server. We are the one company that is doing [Linux] on the desktop and the server across the board, big-time." Then just two days later this view had changed to saying that there was no good reason to buy Linux in the enterprise, and that Linux is a "great environment for the hobbyist".

Since then, Sun seems to have finalised its opinion of Linux, as its commitment to JDS shows, but the prevarication almost certainly cost Sun

quite a lot of goodwill in the community. Furthermore, soon after news of the SCO case against IBM came out, John Loiacono, vice-president of Sun's Operating Platforms Group, said "We bought our Unix license out...we are unencumbered for all things," followed by "For people looking at the issues at hand, we are a safe harbour. We have absolute rights to our technology."

Sun can't really afford to make any more U-turns about Linux, so this is the big push for its move into Linux, and as such can't really afford it to fail. Sun's primary business model lies in producing and selling high-end server machines at premium prices, and during the dot.com boom, Sun experienced brisk sales in these machines. Since then, though, from a stratospheric high of over US\$60 a share, Sun's shares have gone a long

way downhill to just over \$4 a share – along with most other tech companies.

If Sun can make JDS work – and we think it can – it will give the company a viable market share in the desktop computing world, and also, by extension, a larger market share in the server arena. Although companies like SGI, IBM, and HP do a lot to help further the Linux cause, those companies' motives are invariably geared towards shifting more hardware (and you can't blame them); and Sun is no different – if companies are deploying JDS and JES to run their systems, the chances are they will be very eager to eliminate one more point of contact and purchase Sun hardware to power it all.

What this means is that if Sun can really make the Java System work, they will not only give them a slice of the very

lucrative desktop pie, but it will also bolster its lacklustre server sales and give Sun a fresh breath of life.

The encouraging news for Sun is that the first big deployment of JDS has already been agreed – the Chinese government has pledged to deploy one million desktops running JDS over the next year (see *News*, page 10 this issue) with the potential for many more. When Scott McNealy announced this at Comdex, he said "This I believe makes us the number one Linux desktop play on the planet. That's not the only opportunity. We're calling on every ministry of information technology on the planet!"

This is quite a victory for Sun, and proves that JDS is viable even on a very large scale. Now all that remains is for Sun to capitalise on this early win, and turn it into more opportunities.

OTHER BLENDS

With such a large investment in its Solaris operating system, it's not surprising that Sun is planning to release the Java Desktop System on Solaris as well as Linux. We caught up with Simon Phipps, a man Sun aptly names as its Chief Technology Evangelist, at the UK Linux Expo 2003 at Olympia, London, who said that Sun's choice to base JDS on SUSE was largely arbitrary, and that a lot of work had been done to abstract JDS from the OS so that if they ever

want to change the distro at the core it would be a fairly easy task.

Although this distro-agnosticism is all well and good, it also means that it's quite easy for Sun to pick up the elements of JDS and move them to another OS entirely with little hassle, which means a port to Solaris shouldn't be too much work at all. Sun has roadmapped a summer 2004 release of JDS for Solaris, but we wouldn't be surprised if it were released earlier,

particularly given that Solaris already has working and tested ports of GNOME and *StarOffice*. The key problem here lies in the fact that Sun aims to make the Linux and Solaris versions as functionally similar as possible, which will require no small measure of programming magic on Sun's behalf.

When it comes to other architectures, the plans are a little more uncertain. At Comdex, Sun announced a strategic alliance with AMD to build and deploy

Opteron-based systems, and as part of that they agreed to port Java Enterprise System to the Opteron platform. This is certainly a step forward for both companies, and is likely to prove a very fruitful relationship in the long-term. At this time, though, Sun appears to have no plans to port the Java Desktop System to the AMD64 architecture – it will be interesting to see how this pans out after the Solaris release is completed.

critical problems found therein will lie unfixed next to a notice that *Windows Server 2003* is available for a fee. On the flip-side, the Sun-using company would have received software upgrades for those five years, and would therefore be bang up-to-date technologically.

YaST is used as the configuration manager, as you'd expect for a SUSE-based distro, and Sun has simply added its own logo to the top of various dialogs to help project a more unified branding image. This comes

across particularly obviously in places where other brands have been forcibly removed, so you'll find Mozilla working in all its usual form and function, with the exception of it's now being rather unimaginatively called *Web Browser*.




Fresh or instant?

Once you take away Sun's smooth new interface, JDS is based on SUSE Linux Desktop, and comes with GNOME 2.2, *StarOffice 7*, *Evolution 1.4*, and *Mozilla 1.4*. Despite the amount of software, the system requirements are still fairly low – a 266MHz PC with 128MB RAM at the minimum, and a 600MHz PC with 256MB RAM recommended.



“We love Linux, and I hope there isn't any doubt about it”

SCOTT MCNEALY, CEO, SUN, FEBRUARY 2002



***StarOffice 7* inevitably steals the show, and is definitely set to impress everyone who gives JDS a passing try.**

JAVA DESKTOP

The suit is still not preferred attire for the real movers and shakers of the computer world, as shown here by Scott McNealy and AMD CEO Hector Ruiz. Is McNealy apologising for the lack of a plan for JDS to support non-64-bit AMD architecture? Another bone of contention might be McNealy's belief that thin rays are a more efficient computing resource, with a single server processor supporting 20 to 200 users, as opposed to the present vogue for 'space heater' PCs each with its own processor...



The installation, done by *YaST*, was as simple as you'd expect from SUSE's usual releases, and works in precisely the same way except you get a wholly GNOME-based system by default and, of course, there's copious amounts of Sun branding. It's important to note that despite being based on SUSE Linux Desktop, JDS doesn't contain the same package list – some of the packages are those from SUSE Linux 8.2, and others are those that featured in SUSE Linux 9.0.

Having said that, there are some elements of SUSE that Sun hasn't

removed – either because it didn't spot them, or simply because Sun liked the idea. SUSE's multi-language welcome screen remains (also branded Sun, of course), as does *YaST*, which comes preconfigured to get system updates direct from Sun. One curiosity lies in the help system, of which one section is entitled "Java Desktop System 2003 Quick Start User Guide" – perhaps this means the next release will be branded 2004, thereby keeping the JDS brand unversioned.


Unification

Furthermore, the Java Runtime Environment bundled includes the new GTK+ look and feel, which means that all the Java-based applications deployed will blend seamlessly with the rest of the JDS desktop, further

unifying the suite. At this time it's hard to tell quite how well this will work, as JDS only ships with four Java apps by default, of which one is a dictionary and another is *jEdit*. What's really missing is some online tool to help people locate other Java applications that would fulfil their needs, either freeware or commercial, as this would help strengthen Java's presence on the programs menu. As it is, Sun hides the Java applications away in their own menu, under Launch/Extras/Java Applications, and we feel that this explicit separation of Java apps could probably harm the image Sun is aiming to achieve.

It's clear that Sun has gone to a lot of effort to make JDS easy to learn for Windows users – the program start menu, cunningly called 'Launch', is

YaST had had 'Sun' stamped all over it, but its core functionality remains intact.



YOUR NEXT CUP

This first release of Java Desktop System is just that, a first release. Despite having gone through at least seven beta releases and an extensive testing program before release, there are still some features that didn't make it into the final build due to a time constraint. Added to that, Sun is likely to get a lot of feedback from customers over the next few months as the

product is deployed, and so they have already provisionally set a date for JDS 2, and that's Q1/Q2 2004.

What we can expect from the next release is as yet unknown, although as Sun has managed to accomplish an awful lot in the first release despite it only being in development for a little over six months, we can certainly expect great things for v2.



Even the loading screens look good in JDS, proving that perfectly anti-aliased fonts are not just the preserve of word-processing programs.


much shorter than most other Linux distros, and the first four icons on the desktop are by default 'This Computer', 'Documents', 'Network Places', and 'Trash'. Furthermore, few people will be confused at the login screen – its soothing, blue colour and tidily arranged set of buttons make your choices quite clear.

Finally, Sun has entirely done away with the grey-on-black, message-filled boot up and shutdown sequences we're all used to, and replaced it with a simple progress bar and a brief message explaining what's going on. This has been attempted by other

vendors in the past, but few apart from Apple's OS X have managed to get it looking quite so good as this, which provides another strong incentive to make the switch. One curiosity during GNOME startup is the little icon in the bottom-right corner, "Macromedia Flash enabled" – we've haven't the first clue what this means, but presumably it's saying that JDS comes with the *Mozilla* Flash plugin by default. If that *is* what it means, we wonder why it's so important to Sun!

All this combined makes JDS a firm contender in the desktop market, particularly when you consider the

'SUSE and Evolution form core parts of JDS, so do Novell's recent acquisitions seem to be more than just coincidence?'



Die-hard SUSE fans will be pleased to see that the *YaST* ugly installation manager remains almost entirely intact in JDS. Ugh!



"Welcome", "Bienvenido", "Willkommen", "Bienvenue" – what is this, an airline sickbag?! We get the point, boot up the desktop already!

TOO HOT TO HANDLE?

Soon after the *Project Mad Hatter* rumours started circulating, Novell made the first of its two surprising purchases by snapping up Ximian, creator of *Evolution* and *Mono*. Then in November, Novell moved forward again and announced its intention to purchase SUSE. Taken individually, these seem like smart business moves, but if you consider that SUSE and *Evolution* form core parts of the Java Desktop System, it does seem more than just a coincidence?

Whether or not this was intentional by Novell may never be known, but it does leave you wondering where *Mono* is headed with relation to JDS. If Java was Sherlock Holmes, .NET and *Mono* would be Professor James Moriarty, but given that Novell fully supports continuing the *Mono* project, it does make you wonder whether future versions of Java Desktop System will include .NET code somewhere in there – whether Sun likes it or not.

easy and low-cost licensing as well as the promised support plan. Sun doesn't have a great deal of experience in the user-friendly desktop arena so far, but where the company does have a lot of experience is in the field of making customers satisfied, and Sun is certainly pulling out all the stops to make that count for all it can. [LXF](#)



CORRUPT CDS

The rock-'n'-roll spirit of rebellion is far from dead, despite all the awful manufactured pap clogging the commercial airwaves. **Maurice Kelly** mingles with the growing crowd of naughty 'buy, burn, return' teens outside the record shop to bring us the very latest from the hit parade.

A new musical revolution has begun. It has been slowly gaining momentum over the last couple of years and has gone mainstream, turning up on the shelves of even the most unadventurous of music shops. But it isn't what you might think it is. This isn't a new musical style and it's not a musical subculture. It's all to do with the way music is delivered to you.

The Past

First, a detour to the mists of time that was the Eighties, when the recording industry introduced us to those now-ubiquitous pieces of plastic called Compact Discs (CDs) and we lapped them up. They were small, sexy and – according to the BBC's *Tomorrow's World* at the time – lasted forever and will play even after a liberal coating of marmalade. And in recent years, one of the best things about CDs was that they could easily be ripped to your computer and encoded as sound files. Sure, you could make copies of LPs to tape or computer, but it wasn't a digital copy, so a certain amount of analogue noise was introduced.

Then, a number of computer revolutions happened and it all went ballistic. MP3 encoding of sound files meant that they took up considerably less disk space – and more importantly bandwidth – which allowed people to share music more easily. Domestic bandwidth increased as modem manufacturers squeezed more out of conventional modems, then broadband kicked off. And, of course, the advent of peer-to-peer

(P2P) networks made it so much easier for people to share information.

Interesting times were afoot and music-sharing became a nagging headache for the recording industry. It really isn't surprising – with the ready availability of high quality songs at no cost, many people were downloading in their droves and making their own CDs rather than buying them in shops.

In short, the record companies had lost control of their product – from their point of view, it had to be stopped and fast. So they went after the networks that allowed people to share files – the highest profile case being Napster. More recently they've been targeting the people who offer massive amounts of songs to share on present P2P networks. It was a situation like trying to mend a broken pipe while the water is still gushing out. The sensible way to fix it is to shut off the water at the source – and this is where the latest musical revolution comes in: the shut-off valve.

The Present

Like many people worldwide, I am a music collector and own upwards of 500 CDs, records and tapes, and I dread to think how much money I've spent on it all. I listen while I work, while I'm driving, while I relax, and I'm listening while I write this article. I have so many CDs that I created myself a custom MP3 jukebox (see *LXF44*) to save me changing CDs so often. And it was all working gloriously until about a month ago when one of my favourite bands got into the latest musical revolution.

The latest CD by this group has been 'copy-protected' by their record

label as part of a new policy that sees all future releases being 'protected'. Unfortunately their 'protection' of this album, means that I can't produce sound files for listening on my jukebox. I can't listen to the CD in work on my Windows 2000 laptop unless I settle for a dodgy player included on the CD – which plays a low bitrate version. I can't listen to the CD on my Linux boxes at all because the aforementioned dodgy player is for Windows-only. I can't listen to this type of CD in some cars because they contain CD-ROM technology.

So, by 'copy-protecting' the CD, the record companies are trying to shut off the water mains – stop people from ripping and copying, and then they can't share it. It's a short-sighted but obvious solution really. By this stage, you probably have two questions: what is a corrupt CD; and why the proliferation of quotation marks above? The short answer is that many of the so-called copy-protected CDs are in fact violating the specification for CDs and technically are corrupt. But now that that the point is made, I'll drop the quotes and use the terms 'copy-protected' and 'corrupt' interchangeably from here on.

The Effects

Many people won't even notice the difference as they just play their CDs on regular CD players. The people who will notice the difference are those who are a bit more adventurous – who throw their CDs into computers, portable CD players, DVD players, games consoles, and car CD

players, devices that have all had problems with corrupt CDs. Some have reported problems with playback in regular CD players, which means that not only are these CDs preventing copying, they're

preventing playback of the original disc. Unfortunately it seems that the record companies are unconcerned that not everyone can enjoy the CDs they have legally purchased as long as the majority of listeners can.

The methods employed in copy-protecting a CD vary depending on the technology being used. A couple of popular techniques include corruption of the CD's table of contents (TOC) and the introduction of defects to the actual music itself which manifests as pops and skips in the ripped music. For more information, see the box *Copy-protection Schemes* overpage.

Probably the biggest problem is the introduction of deliberate defects into the audio portion of the CD. Home CD players are designed with error correction built-in, allowing them to deal with small errors and continue playing the CD with minimal interruption to the audio stream. This was always one of the big selling points of CDs as you could, in theory, abuse your CD and it would still play (anyone remember that *Tomorrow's World* also kicked their CD around the floor back in the early Eighties? And it *still* played!)

As you'll have experienced, in reality CDs are not indestructible – they pick up scratches and dirt just like everything else, and there's only so much dirt and damage that your player can account for. Over time, a pristine CD becomes harder for your CD player to play back, and it has to do a bit more work to make sure you don't notice the difference. With deliberate defects introduced, your player will see your copy-protected disc as one that has picked up a bit of damage over time, despite it being brand new.

In the short-term, it's not a problem – your CD will play fine. Over time, there is a chance that the deliberate defects will combine with natural ones, resulting in a CD that skips prematurely. A knock-on effect is that your poor little CD player is doing a lot more work than it really should be and, in theory at least, this isn't likely to do it much good at all. While these effects haven't yet been proven with regard to home audio CD players, there have been numerous

ISN'T SHARING MUSIC WRONG?

Groove carefully now, pop-pickers!

The legalities of what you can and can't do with music is a grey area, and gets more confused as time goes on. There can be no dispute that – in the UK at least – you're breaking copyright laws if you copy CDs for friends, or if you distribute MP3s/OGGs of copyrighted material. Yet despite the millions of pounds spent on advertising and radio/television airplay every year, there is one type of free promotion that the recording industry can't really measure but can't afford to take for granted – illegal sharing of music.

Many people have grown up in a culture of sharing music – not 'nasty and malicious' sharing via MP3s, but sharing through the passing around of tapes and CDs and listening to music together. While it is illegal, people copy and share tapes and CDs every day, and people hear things they may not otherwise get the opportunity to do so – and often go on to buy things that

they've heard as a copy. It seems like a natural extension to share audio files with friends, to spread the word and the experience, but it is the wholesale theft of music in digital form that has brought down the wrath of the RIAA.

It is useful to remember that the artists don't get anywhere near as much money as the record companies would have you (and probably the artists themselves) believe. Increasing numbers of musicians, concerned that they're being ripped off by the same record companies they believe are also cheating the CD-buying public, now bypass traditional record companies completely, distributing their music via the Internet for a fraction of the price charged by retail chains: see Ani DiFranco's www.righteousbabe.com/ and www.davidrovics.com/ for a couple of great examples.

Stealing music isn't the answer to a fairer deal for artists or consumers.



If you like Michael Moore's books, you'll love David Rovics's music! With the suspension of mp3.com by the time you read this, his ability to host free downloads of his songs will be severely limited: can any readers host some of his material? Contact him through his site.

reported cases of corrupt CDs causing problems with computer equipment.

Probably the best-known example is the Celine Dion CD that was reported to cause crashes for PC users, and complete lock-ups in some Apple computers. It has now been immortalised in Apple's *Knowledge Base* where the fixes range from software upgrades to taking your Mac into your local service centre just to get the CD back out of the drive. Celine Dion is not alone in causing hardware problems – Mike Oldfield's *Tubular Bells 2003* has also caused some reported damage to computer CD/DVD-ROM drives. This problem is



CORRUPT CDS



not going to go away and it is likely that more and more older equipment will be unable to cope with the increasing levels of 'protection' served up by the record labels.

What About Us?

Unless you're a Windows (or in some cases a Mac) user, then you're not really allowed to listen to your newly purchased copy-protected CD on your computer. OK, that's not true but if you pick up a copy-protected CD from the shelf of your local CD store, you could be forgiven for thinking it. Companies such as Macrovision have gone for the lowest common denominator approach when it comes to making their protected discs playable on computer equipment.

When emailed on the subject by us, the response from Macrovision was that the record companies do not ask for Linux support and so they do not deliver it. Mailing two UK record companies asking about Linux support was also a predictably fruitless activity. Macrovision's stance was justifiable – after all, Windows is the most common operating system out there so there is no point in adding extra features that the record companies are not demanding.

Unfortunately the shenanigans with copy-protected CDs can only get worse for Linux users in the short-term. Copy-protection is a constant race involving content producers and would-be pirates. As a result, the methods employed in copy-protection are becoming increasingly advanced, and so the ability of the average Linux user on the street to play a copy-protected CD on their computer decreases. With more people picking up Linux and giving it a try, the problem is going to become more widespread, and with recent protection systems utilising Digital Rights Management (DRM) solutions, it won't be long before evasion tactics become too convoluted and many will simply give up on listening to music on Linux, which doesn't reflect well. It's bad for the users and it's bad for the image of Linux as a serious contender for the home and desktop market.

This whole thing should feel very familiar to Linux users – remember the heady days of the Content Scrambling System (CSS) and the problems it caused people who wanted to watch DVDs on their machines? CSS was used by film studios to control where and how

COPY-PROTECTION SCHEMES

What protection schemes to watch out for – and avoid!

CACTUS DATA SHIELD (CDS)

The CDS family of protection schemes were originally developed by a company called Midbar Tech Ltd, which was acquired by Macrovision (better known for the technology used to prevent the copying of DVDs to VHS.) They range from CDS-100 (utilised on Natalie Imbruglia's *White Lillies Island* – the first copy-protected CD in the UK) to the latest CDS-300 which allows for more advanced DRM-based solutions, although on Windows only. Probably the best known member of the family is CDS-200 which is on many of the copy-protected discs in the UK today.

The speciality of CDS-200 is the two-pronged approach – firstly they use a multi-session CD with a corrupt TOC in the multimedia session. This means that if you place the CD into your computer it gets a distorted view of the CD, which makes ripping a problem (most CD/DVD-ROM drives read the last TOC on the disc, whereas audio CD players only know to look for it at the beginning.) The multimedia session can contain low-quality, encrypted versions of the tracks on the audio portion, although this session only supports Windows PCs and Macs.

Some people are lucky enough to own drives that read the first TOC (and some people are smart enough to 'disable' the corrupt TOC with a felt-tip pen) so CDS-200 has another trick up its sleeve. Deliberate defects are introduced into the audio portion of the disc

which only manifest themselves when being ripped or copied. Of course, some computer CD drives are better than others and can manage to extract almost perfect audio with a minimum of nasty noises despite the nefarious efforts of CDS-200.

CDS seems to be the most prevalent protection system in the UK at moment (the ukcdr.org list of known bad CDs reports almost six times more CDs using CDS than the next system.) It seems to be evolving fairly rapidly and poses quite a problem for the Linux user.

KEY2AUDIO

It's ironic really – a company that provides equipment supporting MP3s also produces CDs which can't be ripped. Sony's DADC facility in Austria has been developing Sony's own in-house protection system for quite a while now, and it is found on many CDs in the form of *key2audio*. This scheme is known to have used similar techniques to CDS by utilising a corrupt TOC on the CD intended to confuse the reading device. Reports vary from disc to disc as to whether or not the audio data is actually corrupted – as with the CDS family, the levels of protection vary from release to release. It was *key2audio* that caused the widespread problems between Celine Dion's *A New Day Has Come* and iMac computers. It has long been said that Apple users have decent taste...!


More recent Sony releases are beginning to feature the latest protection scheme from DADC – *key2audioXS*. This newer format offers similar facilities to CDS-200 (and

higher) schemes through the provision of encrypted audio files which can be played on Windows PCs, and are controlled by a DRM solution. Of course, if you don't use Windows, then you don't count.

MEDIACLOQ AND MEDIAMAX

If you've read anything recently about copy-protected CDs, then you should certainly have heard of SunnComm's *MediaMax 3* product – an acclaimed (by SunnComm) protection system which relied on the installation of new drivers on a Windows system whenever a protected CD was inserted. These new drivers had the ability to detect a *MediaMax*-protected disc and interfere with audio reads from such discs. The fundamental flaw was that the user only had to disable the auto-running of CDs on Windows and the disc could be ripped perfectly. Amazingly, this is one area where Linux users win out – no drivers for Linux means that such CDs could easily be ripped, as the audio portion of the disc contained no defects.

Before such a poor judgement was taken in releasing *MediaMax 3*, SunnComm were the people behind *MediaCloQ* which was on the first protected CD released in the United States. *MediaCloQ* drew criticism as the discs could not be played on PCs and required the user to sign onto a website in order to download playable audio files – a solution not appreciated by dial-up users; and people who weren't too keen on giving out personal details just to listen to the music they had paid for!



An alternative to copy-protection is the use of exclusive downloads that are only available to people with original CDs.

DVDs could be played, and they didn't seem to care much for Linux users – just like the current copy-protection systems for audio CDs. While the technology is not exactly the same, it certainly isn't worlds apart, and corrupt CDs are likely to become a big deal for the Linux world in the same way that CSS did back in 1999.

Can Copy-Protection Work?

If you believe the companies like Macrovision and SunnComm (that make the protection schemes) then these systems work, and they are here for good. Of course these companies have a vested interest in the need for copy-protection, and it is not impossible that they might make exaggerated claims regarding the need for, and effectiveness of, the schemes they are trying to sell.

Despite the 'advances' being made, the existing copy-protection schemes are not impossible to break. More tolerant hardware can do a much better job of ripping corrupt CDs, and in the end, it only takes one person to upload an album to a file-sharing network and the copy-protection may as well not exist. Any protection scheme is only as strong as its weakest link.

Amazingly, one of the weakest links in the copy-protection debacle is the variation in CDs between territories. It is not unusual for a CD from the USA to be unprotected, when the British equivalent is protected. A recent example of this is the special edition release of *Sleeping With Ghosts* by Placebo – copies from mainland Europe were corrupt, yet in British and

Irish shops it was possible to find exactly the same CD with *no protection whatsoever*.

Let's not forget the similarity to the problems Linux users had with DVDs – numerous groups worked very hard to break CSS, and we have a legacy of projects available now which utilise *DeCSS* and other efforts. If there is one thing we should remember about Linux users is that they don't like being dictated to, and there are a lot of people (and not just Linux people) out there who thrive on breaking copy-protection systems for all manner of devices.

The race for ever better protection can harm only one group of people – the consumers. The sooner the record companies realise this, the sooner they can stop alienating their customers. After all, if they make CDs too hard to play, why should we buy them at all, and what will the recording industry do then?

What Can I Do?

The most important thing that any of us can do is remember that we have rights as consumers. If you purchase a CD which is copy-protected and have problems playing it, then please take or send it back. It is worth remembering that smaller retailers are often simply caught in the middle. It isn't always easy to return a CD if you have already opened it, and especially if it is clearly marked as copy-protected. It seems harsh, but one of the best things that we can do is simply refuse to buy corrupt CDs. Also, make a point of supporting local independent music

stores whenever you can! Keep an eye on the list of known bad CDs on the Web site of the Campaign for Digital Rights (CDR) – it's sometimes possible to get tips on alternate sources of the same CD which are unprotected.

While focussed on the evils of the Recording Industry Association of America (RIAA) the RIAA Radar can be useful in picking out the releases by major record labels in the US. Copy-protection systems cost money and smaller record labels are less likely to release CDs containing protection measures. Pick your future purchases carefully – record companies spend a lot of time analysing sales figures and avoiding copy-protected discs is a good way to get your point across.

If you feel pro-active, you should consider popping along to the CDR site where you can read a lot more about this issue (and others which affect your digital rights.) Read the information, form your own conclusions and, if you agree, tell someone about it. Look out for future leafletting campaigns, and take part if you can – the best way to stop the spread of corrupt CDs is to let people know what they are getting.

In Conclusion

As Linux users, we have a vested interest in eradicating corrupt CDs. These things simply make Linux users the second-class computing people

“If you purchase a CD that is copy-protected, please take it back... and pick your future music purchases carefully.”

that we have fought so hard not to be. The inability to play these discs is a setback to the work done by the coders everywhere who have worked to make Linux a viable operating system for the man/woman on the street.

It is important to remember that your discontent with a corrupt CD can only be seen if you do something about it. Take it back, tell people about it, email the record companies, and avoid them from now on. Start the next musical revolution by saying “No!” to the current one. [LXF](#)

RESOURCES

Here is a selection of links that can give you lots more info about the issues surrounding filesharing, CD (and DVD) copying and use with Linux. Why not head over to the *LXF* forums to discuss the issue of corrupt CDs and exchange useful links with other Linux users?

Campaign for Digital Rights
<http://www.ukcdr.org/>

IFPI
<http://www.ifpi.org/>

British Phonographic Industry
<http://www.bpi.co.uk/>

RIAA Radar
<http://www.magnetbox.com/riaa/>

Tubular Bells 2003
<http://www.rcarter.34sp.com/oldfield/tubularbells2003.html>

Celine Dion vs iMac
<http://docs.info.apple.com/article.html?artnum=106882>

Electronic Frontier Foundation
www.eff.org

The OpenDVD organisation
www.opendvd.org/

What on Earth is... THE AFFS?

Freedom's just a dusty road, heading to an esoteric article about licensing, or something like that. **Richard Smedley** introduces a membership organisation that treasures the freedom of GNU/Linux.

» AFFS – you Linux chaps love acronyms, don't you?

Well, *Association For Free Software* is a bit of a mouthful, but you can pronounce AFFS as a word: to rhyme with the market trader's abbreviation of Wordsworth's favourite blooms – "daffs" – if you like.

» And what does AFFS do?

The Association For Free Software (AFFS) is a membership organisation that promotes and defends Free Software in the UK.

» But isn't that what the Free Software Foundation was set up to do 20 or so years ago?

Well, partly, but not for the UK. The Free Software Foundation (FSF – and there's plenty more alphabet soup to come, I'm afraid) is an international organisation based very firmly in the USA, and closely

involved with the GNU project. AFFS is an associate organisation of FSF Europe, and maintains links with a variety of organisations – collaborating where suitable on a number of projects.

The FSF only recently started a membership scheme. AFFS is a membership organisation, which means that it represents the Free Software community in the UK.

» All of the community? I thought that programmers couldn't agree on anything?

Well, all those that decide to cough up the modest membership fee. Until the founding of AFFS (covered in *LXF26*, pages 9 and 20), Free Software lacked a campaigning voice in the UK, and none of the existing professional or voluntary organisations fitted the bill. Similar organisations exist in most other European countries. Of course a range of views is represented, which leads to the occasional

lively debate. However as the organisation exists solely to promote and defend Free Software there is no disagreement on the fundamentals.

Oh and it's not just programmers – it's users too. Which really means most of the population are potential members, so join now and beat the rush!

» OK, so how does the AFFS "promote and defend" Free Software in the UK?

Promotion takes many forms, from distributing funds, to talking to schools and the government. Running conferences and appearing at talks and exhibitions. Advertising, producing leaflets and articles, it's a busy organisation.

Many AFFS members are quietly working away at promoting Free Software in their jobs and other areas of their lives, but the strength of the organisation is that it brings people together of diverse skills and talents, to make the most of campaigning opportunities.



»» And defending – does Free Software really need defending?

As for the defensive part, that's a matter of planning for and meeting threats that come mostly from legislation – such as the current software patents brouhaha. Legislation over copyright can also cause problems for coders reverse-engineering programs to ensure file-compatibility with market leaders who have little regard for open standards.

As well as countering legislative threats, AFFS also campaigns to keep a level playing field for companies tendering for government-funded projects. Recently, committee members were involved in government consultation on the use of Free Software by the government – government purchase of IT has been under fire from many quarters.

»» But won't patents protect small software houses (including those that produce Free Software) from large corporations muscling in on their innovations?

Short answer: No. Software is adequately protected by existing copyright legislation.

The long answer would need an entire *What On Earth...* feature, software patents harm innovation and only benefit the largest companies who hold a raft of patents and, of course, the lawyers – for more details see our article *Freedom To Innovate* in *Linux Pro*, August 2002 (free with LXF30) – or www.cl.cam.ac.uk/~mgk25/stallman-patents.html

Recently criticism of patents on software has been growing in the USA too with the US Federal Trade Commission calling for a change, details at www.theregister.com/content/4/33691.html, and there have been calls for an overhaul in the law from Intel co-founder and chairman Andrew S Grove, who warns that the US is losing its competitive edge in key tech sectors (<http://swpat.ffii.org/news/recent/index.en.html#washingtonpost031010>). The European Parliament overturned the attempt to introduce software patents here, but at the time of writing, it is with individual member states to overturn the European Parliament's decision. AFFS will be watching developments closely.

»» Hang on, did you say distributing funds?

There are many small (and even a few not-so-small) pieces of important software unavailable under free licensing terms – some of them are quite UK-specific, such as localisation of accounting software for UK tax regulations. A little money can go a long way in pushing along an important program.

»» So where do you get all this money – I thought that you said that membership fees were modest?

They are indeed – at £10 per annum they have been set so that they are not a barrier to joining.

These fees, together with donations and sponsorship, cover running costs – manning shows, printing leaflets, etc. Most of the money for UK Free Software projects is generated by UK Free Software Network, an ISP which generates money for Free Software – AFFS is the third-party that makes the decisions over who gets the cash.

»» Can I have some?

Tell us about your project. Why is it important? We are open to applications.


Twenty years ago, the UK government paid British programmers to write software for British schools on the BBC model B microcomputer. Now millions of pounds of taxpayer's money goes instead on licensing fees for proprietary software from an overseas company, whose products we have to use 'as is' with no possibility of modification.

In Germany, the government recently funded *Kroupware*, as they saw the need for better groupware on GNU/Linux. Until the UK government sponsors local coders to fulfil local needs, AFFS will do what it can to fill the gap. In fact the government has recently funded the APLAWS project to provide Free Software to run local authority websites, however it requires some non-free software to be used. AFFS will fund projects for the benefit of the community, not just to meet government targets.

»» You talk a lot here about the government, what about AFFS and business?

Business has been much quicker to realise the benefits of Free Software when it comes to control, economies of scale and avoiding vendor lock-in. Governments are slow to change and sometimes lack critical technical expertise, preferring to listen to a few 'industry insiders'. However, the unstoppable success of GNU/Linux around the world has seen many countries adopting it at a national level. For this to happen here, your help is needed to add your voice to the campaign. You will be aware from LXF's news pages that Free Software is mandated for government use in parts of South America and across the two-thirds world. In Europe, France and Italy have made a number of moves towards promoting Free Software use, particularly in education. Spanish regions such as Extremadura have more wholeheartedly embraced GNU, sending out localised install disks to the local population like some fervent Free Software AOL. German cities find themselves in the news for ignoring Microsoft's lowest bid.

Governments spend a lot of money on software and put it in front of a lot of people – a small part of this reaching Free Software projects will be a real boost to the platform. As to the more conservative small- and medium-sized enterprises (SMEs), AFFS is moving towards creating resources to help them find Free Software solutions, and to grow



WHAT ON EARTH AFFS

« the market for Free Software businesses in the UK – many of whom are supporters of AFFS. Again this is an area where readers of *LXF* can help to swell the campaign, and talk to local business organisations. Indeed some notable LUGs already do just this.

»» This all sounds interesting, but just how accountable is this group?

AFFS has a committee charged with day-to-day running of the organisation. Elections for committee posts take place at the AGM. The constitution is being fine-tuned to ensure the widest possible participation in the elections by the members. Accounts are published regularly and transparency is valued highly – it's in the nature of the Free Software community.

»» Very laudable, but surely not many people want to attend a boring old AGM?

Probably not, which is why the AFFS Annual Conference (AFFSAC) runs on the same day. The inaugural event, in 2003, was a great success in spite of the sudden change of venue following the postponement of the *LinuxUser* expo. There were many interesting talks on Free Software projects such as the *Rosegarden* musical composition software, *Moodle* content management system, and the *Kyfieithu* project which is localising KDE (and GNOME) for the Welsh language. Members travelled to Birmingham from all over the UK, and many more joined AFFS specifically to be able to attend.

For 2004, a bigger event will be announced soon, when a venue has been fully secured. Watch the website (or the announcement mailing list) for details.

»» So you really want the membership involved?

AFFS has moved its main projects to a workgroup structure to give those most interested in a particular area, such as legislation, more control over achieving AFFS's aims. New ideas for workgroups can be submitted to the committee via the website. AFFS will support them with server resources, advice, press releases, newsletter space, adverts, organising donations and other help.

»» Government, business, community – you seem to have it all covered.

AFFS is actually just beginning – it could actually be doing a lot more – with your help! Every member adds to the weight of its campaigns. A minister is much more likely to listen to an organisation with 100,000 members than one with 100. And of course the campaigns need as much involvement as you can give – whether it's talking to your MP, writing an article for your local newspaper or just burning a couple of Knoppix CDs for your children's school or local library. And of course there's much more to be done in talking to schools, who are turning out the next generation of Free Software users.

Education embodies many of the values of Free



Software – principles of sharing and community – but is also important as schools are very concerned with social inclusion, and the fact that schools can give the software needed to pupils to run on hardware which is nowadays finishing up in skips, means that Free Software is an obvious answer to bridging the 'digital divide'.

»» Education, education, education?

It's very important – taxpayers spend tens of millions of pounds annually on software licences for schools to use programs that are often used for training rather than education. This winter AFFS is jointly sponsoring a conference in London for those involved in bringing Free, Libre and Open Source Software (FLOSS) into education. The keynote speaker will be Professor David Hargreaves, formerly Professor of Education at Cambridge University, Head of the Qualifications and Curriculum Authority (QCA) and currently Chairman of BECTa (British Education and Communications Technology Agency).

The event will take place on 17 February 2004 (during half-term) at the London Institute of Education. Aimed at IT staff in schools, headteachers, IT advisers and inspectors, support companies and specialist maths and computing schools. Case studies will be an important part of the event, and if you know a school about to make important ICT decisions, tell them that they should be attending FLOSSIE.

»» FLOSSIE?

The Free, Libre and Open Source Software In Education conference, yes it's a catchy name. FLOSS is the current term for Free Software that covers the term Open Source, in a bold attempt at inclusiveness.

»» Open Source – I was wondering why you hadn't mentioned that.

Why Free Software and not Open Source?

From the FAQ:

"Short answer: because that was the name the founding members preferred.

"Longer answer, but still incomplete: there are many reasons, including a more concise basic definition, but almost the same set of accepted copyright licences. Many groups calling themselves 'Open Source' declare that benefits which aren't technical or economic are off-topic. Some refuse to condemn use of software patent licences to restrict freedom to use software, while we campaign against software patents. An early intention of OSI was to secure a trademark on the term 'Open Source' but they failed and some think the term is too ambiguous now. Some who use 'Open Source' claim the name doesn't matter anyway, as long as the idea is right..."

In other words, AFFS is unafraid to talk about 'freedom' while still talking to businesses about the benefits of the software to all. While Free Software and Open Source software are very different animals, both camps are currently campaigning in the same

area (replacing proprietary software with GNU/Linux and BSD), and make up the two halves of the 'Linux community' in this country.

» One big happy family. Do you take money from anyone?

AFFS exists to promote Free Software. It will work with any interested party on particular projects, provided there is no conflict of interest.

AFFS is grateful to accept corporate sponsorship for particular projects (Positive Internet provide AFFS's webserver, UKLinux paid for the leaflets distributed at the recent *LinuxUser* Expo), and sponsors are aware that AFFS's acceptance of the money does not in any way endorse the particular company. For many companies working with Free Software, a donation could be seen as an investment in the future.

» You put GNU in front of Linux all the time – why not OOo/KDE/X Windows System/GNU/Linux?

GNU is an operating system, started 20 years ago to replace proprietary UNIX. After ten years of hard work, the FSF's OS lacked only a kernel. Into this void came Linus Torvald's celebrated collaborative effort, and a whole new audience for GNU. OOo, KDE and the X Windows System are not operating systems, however essential they are to many modern users. The title GNU/Linux acknowledges the work of both the Free Software Foundation, in producing the GNU operating system, and the Linux community, in bringing us a robust kernel. *LXF* affords GNU/Linux proper recognition every month on page 114.

Some would say that to call the OS 'Linux' is not just to ignore all the work of the GNU project, but to deprecate the values of freedom that inspired, and still drive, the FSF – which may be unknown entirely to new users coming to a 'Linux' distribution that bundles proprietary software, and speaks of the Open Source advantage. That is why AFFS endorses the FSF-Europe campaign: "We speak about freedom".

» Just why exactly should software be free?

Most software doesn't appear shrink-wrapped in boxes on the shelves of PC World. It is written to order, to solve specific problems, for countless businesses throughout the world. Many of these problems bear strong similarities to each other, and could more easily be solved with reference to existing programs – which could be readily adapted to new situations.

That is they could be readily adapted were they available under a Free Software licence. In no other fields of science do practitioners find themselves constantly 're-inventing the wheel' as they go back and rewrite algorithms that have been written many times before by several programmers before them. What is the benefit of that?

» I thought that I was asking the questions?

Indeed, sorry. But one wonders what would be the state of particle physics if each research team had to rediscover sub-atomic particles, already known to others, before advancing the science with new experiments. However, do ask another question.

» Everybody seems to have their own licence which is like the GPL or like the BSD licence. I'm confused, but diversity is good, right?

It's not the similarities that count for the licences, it's the things that the licence does differently. What is it that the licence forbids you from doing?

The modified BSD licence is a simple permissive licence that allows you to do what you like with the code, including hiding it inside your own proprietary application. Version 2 of the GNU General Public Licence (GNU GPL) is an elegant statement of freedoms designed to protect the rights of the user, by passing on freedoms to each subsequent user of the software.

» So the BSD licence is more free – as it lets you take the code (allowing Microsoft to use FreeBSD's TCP/IP stack in several versions of Windows)?

It lets you take code from the community – it's a freedom to take away from those who write, test and use the code by choosing not to contribute back changes and improvements that you make. The GNU GPL *protects* the freedoms that it gives.

Proponents of both licences are most vociferous in insisting that 'theirs' is the most free licence, and that the other is the work of the devil. Back in the real world, both are Free Software licences and both are acknowledged as such by the Free Software Foundation and by the AFFS.

» Remind me, why does all this matter to me?

As a GNU/Linux user you have a vested interest in the continued availability and success of Free Software. Some companies, including a few with a great deal of money, have interests that would benefit from the failure of FLOSS. As GNU/Linux becomes a serious threat to established software companies, proprietary software vendors will try their best to make conditions tough for Free Software, you can help it to succeed.

» Where do you stand on copyright and music?

We are the Association For Free *Software* – many of our members may (or may not) have radical views on the issue, but there are other places for them to pursue such topics. We find we have more than enough work getting Free Software noticed in the places where it is most needed (government, schools, the health service, charities, small businesses, *et al*), and we certainly don't want to dilute our message by trying to agree a stance on more general issues.

» Can I get on your mailing list?

AFFS has its own lists for internal business and for announcements, to which members subscribe, but a lot of general discussion takes place on the FSFE-UK list, to which anyone can subscribe: <http://mail.gnu.org/mailman/listinfo/fsfe-uk>

» You've convinced me to join – where shall I send my tenner?

You can join online, or print off the form from www.affs.org.uk/documents/membership.pdf If you have yet to pass on this issue's coverdisc to a Linux-curious friend (ah, the benefits of Free Software) there's a form on there too. [LXF](#)

ALPHABET SOUP

Let's expand those acronyms, and add in a few URLs...

AFFS – The Association For Free Software
www.affs.org.uk

Postal address:
c/o Turo Technology LLP
79 Sir Lewis Street
King's Lynn
PE30 2AL

AFFSAC – AFFS Annual Conference
See www.affs.org.uk/affsac.html

AFSP – Association of Free Software Professionals
"promotes the publication of high quality, standards compliant, Free Software and dispels the common misconception that Free Software is of poor quality and inferior to its commercially-produced counterparts."

BSD – Berkeley Standard Distribution
The Unix which forms the basis of FreeBSD, NetBSD and OpenBSD. The modified BSD licence is a model of the *soi-disant* 'permissive' licence.

FSF – The Free Software Foundation
Set up to nurture the GNU project, and the Free Software movement in general. <http://fsf.org/>

FSFE – Free Software Foundation (Europe)
Bridging differences in time zones and culture with the USA to keep the Free Software torch burning here. www.fsfeurope.org

GNU – GNU's Not Unix
This is the project to produce a Free Unix that was started by Richard Stallman in 1984.
www.gnu.org/

GPL – General Public Licence
Sets out four essential freedoms for software users.
www.gnu.org/licenses/gpl.html
Find answers to your questions at
www.gnu.org/licenses/gpl-faq.html

LUG – Linux/Local User Group
Organisation for GNU/Linux users at a local level.
See www.lug.org.uk

RMS – Richard Stallman
<http://stallman.org/>

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

Don't fear the command-line – step outside the GUI and into a whole new world **p62**

HTML/CSS

PART TWO: Using HyperText Markup Language and Cascading Style Sheets with Linux **p66**

Blender >>

Mesh techniques: how to create complex objects by stretching a 'skin' over a framework **p70**

The GIMP

Filter Fantasies Part One: some crafty non-destructive ways of tweaking your digital photos **p74**

Practical PHP

Object-oriented programming, and what the future holds with a first look at PHP 5, plus an introduction to Access Control **p78**



ILLUSTRATION BY SHIGETO MAEDA – www.blender3d.org

Server School

Long before the intricacies of HTTP, FTP ruled the Web – and it's still a great way to transfer large files **p82**

TCP/IP and Linux

Linux is a complete solution, not only for application services, but also running underlying networks **p86**

TIP OF THE MONTH!

BUYING A LINUX BOX – UK

As regular readers will know, *Linux Format* is interested in tracking the experiences of private individuals trying to buy a desktop or laptop that comes pre-loaded with a Linux OS.

Thanks to Geoff Bagley, Jeff Rhys and Mark Jose – all of whom CC'd us in their prospective Linux PC purchase enquiries to customer.care@evesham.com – *Linux Format* will be contacting many computer retailers in 2004 to further investigate why it's so difficult to buy Linux PCs in the UK. Special mention must go to Phill Rogers, whose request for a refund on Microsoft XP because he stated that he didn't agree with the EULA

produced a story (and questionably legal) "You can remove the OS, that's entirely your decision as it's your machine, but we would not offer a refund for it as we do not sell the machine without the OS," reply from a representative at Evesham.

Reader Derrick Simpson wrote telling us about **Play-Net Computer**, 50 Baghill Lane, Pontefract, West Yorkshire WF8 2HB. Tel: 01977 690909; email sales@castle-computer.co.uk. It has for the past two years or so supplied computers with SUSE Linux pre-installed. Play-Net Computer has also been a SUSE reseller since version 6.3, introducing

Linux and Open Source to many clients including SMEs. The lowest price machine aimed squarely at newbies comes in at £399, with a wide choice of options and upgrades. Customers can customise their new PCs from the start or add upgrades later, as cost and experience permits.

If you CC lxf@futurenet.co.uk with all enquiries that you send to any UK retailer or manufacturer about purchasing PCs with Linux (or without an OS), we can monitor the situation more closely. Also tell us about any retailers like www.transtec.co.uk that sell Linux PCs in the UK so we can pass the news on to our readers. **LXF**

BEHIND THE GUI

Beginners' Guide to Linux: The Command Line

At some point the dreaded command line has to be faced and all that graphical loveliness cast aside, if only for a short time. But don't worry! **Andy Channelle** is here to hold your hand...

Earlier this year Dylan Evans, co-author of *Evolutionary Psychology* (ISBN 1-8404-6043-1) – a book which reportedly influenced the structure of *The Matrix* film trilogy – suggested that computer users should look beyond “those pretty little icons” on the computer screen, and achieve a more truthful relationship with technology. What if, he asked in *The Guardian* (www.guardian.co.uk/print/0,3858,4791087-111020,00.html), the flashy graphical user interface (GUI) disappeared and, like Neo in the film, you were left with just ‘the code behind the graphics.’

Well, after so many months skirting around the issue in this series, that is the goal of this tutorial. It is quite a daunting prospect for both author and reader but, as Evans says, “if you want to be truly free, you have no choice but to understand the machines you work with.” Lock and load people!

The Command Line Interface (CLI) is the precursor to – and base of – the Graphical User Interface (GUI). As a simple example, when a user double clicks on an icon to launch *XMMS*, for instance, they are merely executing a command. In this case **xmms**. You could get the same result by opening a terminal, typing **xmms** and hitting the **Enter** key. To see the command, right-click on the *XMMS* icon (or other application on the desktop) and select ‘Properties.’ The command will be listed under the Execute tab.

Obviously, commands can obviously do more than just start applications. In fact, anything you can do with a GUI (and a whole lot more) can be accomplished with the command line. The difficulty is that where the GUI is pretty and mostly intuitive, CLI

work is text-based and relies on hundreds of archaic terms – hence the common perception that the CLI is only for geeks or people with inordinately large bookshelves.

Some basic stuff

The first thing we need is a terminal. Most distributions ship with a wide selection of terminals (or ‘consoles’ as they are sometimes called) and there is usually an icon handily situated on the taskbar or KDE or GNOME that will open the default application.

One of the first things the novice may be expected to do in the daily running of a Linux PC is ‘become root,’ *ie* gain root privileges in order to change some system properties. Remember that the inherent security of Linux relies, for the most part, on the system of allowing users to change only a limited range of features. To really do any real damage, you have to log in as root, hence the constant stream of ‘don't run as root’ advice that you'll hear from most sources of information about Linux – including *LXF*.

However sometimes, for instance when building an application from source or installing something ‘system wide’ you will be called upon to become root. You can gain those privileges by typing **su** and the root password. All commands typed in that terminal will then be recognised by the computer as coming from the system administrator, until you type **exit**.

KDE has its own extended version of **su** which allows, for instance, the editing of system files from within *Kwrite*. Type **kdesu kdwrite**. You will first be prompted for the root password before *Kwrite* loads up with root privileges, meaning you can open and edit system files from within a familiar graphical text processor.

The most basic CLI tools deal with file management, so here's a few to begin with, they all follow the same formula, a command followed by one or more ‘arguments’.

cd – this is a command to change directory, so it needs to be followed by a path, that is the location within the file hierarchy. So **cd /home/andy/Music** will take me to a directory called **Music** which is within **andy**, a sub-directory of **/home**. You will probably have noticed that every user (except root) specified on the login manager will have their own named directory inside **/home**. Simply typing **cd** with no arguments will take a user back to his or her home directory.

Once inside a directory you can list its contents with **ls**. This can also be done using a full path: **ls /home/andy/Documents** displays a long list of files and folders from my Documents folder.

It is possible to copy files from one location to another using the command **cp**. The following text string will copy a file from

RPM ON THE COMMAND LINES

Pick a letter, any letter...

While it is easy to install an RPM by just double-clicking on it in *Konqueror*/*Nautilus*, using the command line actually provides many more options for the management of your Linux box.

The most basic rpm command would be **rpm -i application.rpm**, which simply installs the application in the default location with the default options. However, there is a lot more you can add. Here are the most common options:


- U** upgrade a package.
- h** display progress bar of install.
- e** remove a package.

- q** query a package.
- p** perform the query on an uninstalled package file.
- i** display useful package information.
- l** list files in a package.
- v** display information verbosely.

To upgrade our application we would type **rpm -Uvh application.rpm**. This would also give plenty of feedback about what is going on under the hood and how long it is likely to take.

To remove the application we'd input **rpm -e application.rpm**.

TUTORIAL Beginners' Linux




less allows users to examine and search the contents of a text file.

◀◀ Finally, type **make install** and hit the **Return** key on your keyboard. When the whole process is finished, do **exit** to get out of superuser – root – mode.

Most often, you will find the application in `/usr/local/bin` and if this is in your `PATH` file you should be able to run the application by simply typing its name.


Compiling from source can leave a lot of mess on your hard disk, so it can be useful to follow up a successful install with **make clean**. However, before you do this, *make a backup of the makefile* as it could be useful if you come to uninstall the application later.

Removing self-built software can be easy or very difficult, depending on the work the packager has done for you. The first step is to restore the makefile to its original location and attempt to do **make uninstall**. This will either result in a lot of screen output, which means the system is probably working, or nothing, which means it isn't. If it doesn't work, it is a case of combing through the file structure and manually deleting stuff – a time-



man pages hold a lot of useful info that you need about programs and their processes.

TIPS



When you are working at the command line, you don't necessarily need to keep typing out the same thing. A history of everything you've typed is available by using the **Up** and **Down** arrows to cycle through previously entered commands.

You can narrow down the search for a specific command by typing a few letters of it, for instance **gr**, and then hitting the **Tab** key twice to get a list of every command beginning with those letters.

consuming (and sometimes dangerous) pursuit. In this case, the **makefile** will still be very useful as it will reveal exactly where everything has been secreted.

Useful extras

less

less is a command that lets you view text files. Again it is **command + argument**, so **less readme.txt** will open **readme.txt** in the terminal window, allowing you to navigate a single page at a time. The **Page Up** and **Page Down** keys provide basic navigation options, and you can also type **G** to skim to the end of the file or use **1G** to go to the start. You can search for a specific string in a file by typing **/string** where **string** is what you are searching for, and hitting the **Q** key will take you out of **less**.

file

Often you'll come across a file in the Linux file structure without an extension, and you can use the **file** command to determine the format. Follow the command with a file name to see a short description of the format (for instance ASCII English text, with very long lines). You will then know if the file is viewable with **less**, or if some other viewer/editor is required. **file** can recognise quite a range of file types including `.html`, `.jpg`, `postscript`, `.zip` or `.tar` archives and RPM files. You can examine the contents of RPMs using **rpm -q**.

grep

grep stands for Global Regular Expression Printer. This is a useful tool that searches for specific text strings within a file (or files) without first opening it. So you could type **grep -i linux readme.txt** to output to screen all the lines in the file **readme.txt** that contain the word **Linux**. The **-i** modifier simply tells **grep** to ignore letter case, so we would get all instances of **Linux**, **linux**, **LINUX**, *etc.* It is also possible to print out all the lines of text that do not match a pattern by using **-v** in the command.

head

When you think you've found the text file you want, you can quickly examine the first few lines of it by using the **head** command. By default, typing **head filename.txt** will display the first ten lines of the file. You can alter this by adding **-n** – where **n** is a number – to the command: **head -20 filename.txt**.

tail

This is the same as **head** but displays the last lines of a file rather than the first. It has the same options too.

find OR locate

This works in the same way as most search tools. Typing in

GOOGLE GURU

While we're typing strange things, here is a quick guide to some of the tools hidden away beneath Google's minimalist front page.

If you often search for Linux-related stuff and tire of typing Linux at the start of every search, you can bookmark the site's Linux-centric section – www.google.com/linux – to avoid it. There are also specific engines for BSD, Mac, Microsoft and the US government.

Many people use Google as a glorified spell-checker, but it can do more than simply correct bad spelling. For instance, once your search results are displayed, you will see a blue line near the top of the page with the words: "Searched the web for <search string>". Clicking on any of the words within this line will take you to a definition of that word from dictionary.com. If that's not enough you can use the 'define:' operator to search the web for a range of word definitions. 'define: ambient' will provide results similar to the ones shown in the screenshots here.

The search engine handles numbers too. Enter 2+2*2, hit **Enter** and you'll be presented with the answer – which is six, by the way. In addition to basic mathematical functions, the Google Calculator is pretty good at weights and measures conversions, able to handle miles to kilometres, ounces to kilograms, cups to teaspoons and many others. The syntax is



Searching Google by filetype.

find readme.txt will find all instances of files with that name within the current directory. You can also add 'wildcards', meaning find will search for all files which match a partial name, with the asterisk sign. **find read*** will find all the **readme.txt** files, but could also find **readyou.htm** and **readaloud.exe**, for instance.

gzip/gunzip

These two commands, when followed by a path or file name will compress or uncompress the file using the zip format. **gzip hugespreadsheet.xls** would compress the file (giving it an added .gz extension), while **gunzip hugespreadsheet.xls.gz** would unpack it to the current directory.

man

Last but not least is the **man** (short for manual) command, this

<value> to <value>. If you choose quite a common conversion – **36.3 miles to kilometres** for instance, the resulting page will also contain links to websites that mention the figures.

For those occasions when you know what you're looking for, an OFSTED report say, and don't want to wade through the 1,760 pages that even a county-wide search generates, you can narrow the criteria by defining a file type. In this case, we know OFSTED posts its reports on the Internet using Adobe's Portable Document Format with the .pdf extension. Amending the search to include **filetype:pdf** (with no spaces) excludes everything that isn't in that format. Google currently recognises 12 file types: Adobe PDF, Adobe PostScript (ps), Lotus 1-2-3 (wk1, wk2, wk3, wk4, wk5, wki, wks, wku), Lotus WordPro (lwp), MacWrite (mw), Microsoft Excel (xls), Microsoft PowerPoint (ppt), Microsoft Word (doc), Microsoft Works (wks, wps, wdb), Microsoft Write (wri), Rich Text Format (rtf) and Text (ans, txt).

Finally, you can discover how popular your own website is by commanding Google to check who is linking to you. Simply type **link:websiteURL** without a space to get a list of linkers. Doing **link:www.linuxformat.co.uk** reveals that about 620 other sources link to the *Linux Format* website.

displays the Linux CLI equivalent of the help file in the terminal window. Navigation around the **man** files follows the same scheme as **less** (first point on previous page). Many Open Source applications have their own man pages so it's always worth, in the first instance of trouble, typing **man application_name** to see what is available. And if you really need help with **man**, you can of course type **man man** to find out more.

In fact, you can access the man pages for any of the commands outlined above to discover the variety of options available from the command line. When it comes down to it, working with a command line interface is just a question of either a good memory for commands (which will come with repetition) or access to the right resources. It's nothing to be scared of! [LXF](http://www.linuxformat.co.uk)

NEXT MONTH

Next issue we'll take a look at Sound, including storage and playback options.

In *LXF51* we will start probing the secrets of specific applications. If you have any ideas about what you would like to see covered, visit the *Linux Format* forums at www.linuxformat.co.uk to make your voice heard, or let me know at andychannelle@yahoo.co.uk.

WEB PROGRAMMING

Using HTML and CSS with Linux

PART 2 Jono Bacon continues our quick foray into HTML use by folding CSS into the recipe.

Last month we started our journey into the world of HTML, and we focussed on what exactly it looks like and how the tags work. We also looked at simple examples of HTML tags such as the simple text formatting tags and the table tags. This month we are going to take our knowledge of HTML further and actually use it to create a page design that you can play around with.

Before we get started, it is important that you have the right tools available to make use of the examples that I will be using. You will need a text editor such *Vi*, *Emacs*, *kate*, *gedit* etc (I am using the *Quanta* web editor), and you will need a web browser (I am using *Mozilla Firebird*).

A question of design

Web design can be a very complex subject. If you look around the Internet, there are countless designs and layouts, but most of them follow some common techniques and guidelines, and the vast majority use HTML (some use Flash or Shockwave for their interfaces). The common factors among these sites is the usually the following areas, which should sound familiar to you:

Title The title area of the page shows the name of the website and can be considered the 'branding' of the site. This area is particularly important so people know the name of the site so they can return to it.

Navigation The navigational area of the website is where the user can select which content to view. This area can contain buttons, icons or text to link to these different areas. In most websites, the navigation area always remains in the same place so users always know where to access content.

Content area This is where the main meat of the site is displayed. In this area you will display your article, news, tutorial, information etc. This is usually displayed in the remaining area of the page.

Before you begin creating your pages, it is a good idea to go and view some existing sites and look for these three different areas of each page. You will see that on many sites, these areas are often in the same place on the page.

Starting development

There are two major ways of laying out content on a page: using TABLEs, and using DIVs. The first method is to construct the design using a number of nested tables. The idea is that you create a table with, for example, one column and three rows. Next, you will create a table in the top cell which might have two columns and three rows, and have different table configurations in the other cells. This way you can create a layout of table cells

that are then filled with content. The other method is making use of the DIV tag to create defined areas of content on your page. These DIV tags cordon off a part of the page and allow you put content in there. These DIVs can then be controlled by Cascading StyleSheets (CSS) to define their position and look.

Of the two methods, the better choice is to use DIV tags to construct your layout. Admittedly, nested tables offer a conceptually easier method of creating a complex design, but using DIV tags gives you cleaner code. The other major benefit of DIVs is that you will have accessible code for those using screen readers and other disabled access devices. Many of these devices cannot interpret nested tables, and you want to ensure that as many people as possible can read your site.

I will not be covering the use of nested tables in this article and will instead just look at the better method of using DIVs. Tables should be used for typical tabular information, and we will use them in our layout later.

To get started, create a new file called `bob.html`, and add the following code:

```
<!DOCTYPE HTML PUBLIC "-//W3C//DTD HTML 4.01//EN"
"http://www.w3.org/TR/html4/strict.dtd">
<html>
<head>
<title>Bob's Computer Shop</title>
<style type="text/css">

#title
{
position: relative;
background: #ddddd;
}

#nav
{
position: relative;
background: #bbbbbb;
}

#content
{
position: relative;
background: #999999;
}
</style>
```




Fig1 Simple layout of a page using DIVs.

```

</head>
<body>
<div id="title">
Bob's Computer Shop
</div>
<div id="nav">
<a href="home.html">Home</a> -
<a href="computers.html">Computers</a> -
<a href="services.html">Services</a> -
<a href="contact.html">Contact</a> -
<a href="links.html">Links</a>
</div>
<div id="content">
Welcome to Bob's Computer Shop!
<p>
At my computer shop I can offer a wide variety of services and
products, all backed up with an impressive support plan. We
can build workstations, servers, render farms and clusters
running a variety of Operating Systems including
Microsoft®; Windows®; Linux®; Solaris®; and a
variety of Apple®; MacOSX®; solutions.
</p>
<p>
We can offer a number of pricing solutions including the
following:
</p>
<table border=1>
<tr>
<td>Workstations</td>
<td>Starting at £500</td>
</tr>
<tr>
<td>Server</td>
<td>Starting at £1000</td>
</tr>
<tr>
<td>Render Farm</td>
<td>Starting at £2000</td>
</tr>
</table>
<p>
Please get in touch with any queries. You can email us <a
href="mailto:bob@bobscomputers.com">here</a>.
</p>
</div>

```

```

</body>
</html>

```

The output of this code can be seen in **Fig1**. In this example, we have created three DIV areas. The first is the title area (light grey), and this area contains the code to show Bob's Computer Shop. The second DIV contains a number of links and this is the navigational area (mid-grey). The final, and third DIV is the main content area and contains a few paragraphs and a table (dark grey). Although simple in its design, this example contains the core elements of a page, and we can now begin laying it out in different ways.

Grokking CSS

CSS is essentially a way of styling the contents of your web page. On this page we have a number of elements that we have used (such as paragraphs with `<p>`, tables with `<table>`, `<tr>` and `<td>` etc). Each of these areas can be styled and formatted in painstaking detail, and even at this early stage, we have styled some areas with CSS.

If you look at the top of the code, you will see the following:

```

<style type="text/css">
#title
{
position: relative;
background: #d4d4d4;
}
#nav
{
position: relative;
background: #bbbbbb;
}
#content
{
position: relative;
background: #999999;
}
</style>

```

In the `<style>` tag we have a number of lines of special code that are CSS definitions. The `<style>` tag lets you put in CSS code that will apply to all the elements on the current page. This code is comprised of three blocks – the `#title`, `#nav` and `#content` blocks. These blocks of code are used to format our separate DIVs on our page. Each block title has a `#` before it to specify that it is unique. Using CSS, we can format a number of different tags, and some of these tags can appear more than once on a page. DIVs are slightly different in this regard and generally only appear singularly. In our case we have a single title DIV, single navigation DIV and single content DIV. If you look at the DIV code you will see that each DIV has an `id` attribute that is used to give a unique name to the DIV. Using this `id` attribute, the CSS code above can be applied to it.

To clarify this, we will look at the title DIV. This DIV is given the `id` 'title', and in our CSS block we have a corresponding `#title` block. In this block we have two entries; a position and background definition. These are examples of CSS definitions that will be applied to the title DIV. The position definition allows you specify what type of DIV we are using. This is set to **relative** so that the

GETTING HELP

CSS is a huge standard and comes in two major forms with additional standards in development. The CSS1 and CSS2 definitions should be well established in most modern browsers, and you can read the specification at:

www.w3.org/Style/CSS/

You can also get a lot of help by using the W3Schools website:

www.w3schools.com/

Finally, if you want to see the pure versatility with CSS, have a look at the CSS zen garden:

www.csszengarden.com/

TUTORIAL HTML & CSS

◀◀ DIVs will all flow on the page relative to each other. When we use the **relative** type of position, each DIV will flow into area not taken up by another DIV (such as water flowing into a space not taken up by other water). The other major type of position used is an **absolute** position. This type will require us to specify the exact dimensions and position of the DIV. Although absolute can be more precise, it can cause a few headaches for different sized browser windows.

The other definition used in our **#title** block is **background**. This simply allows us to specify the background colour for our DIV. There are other features of this definition that we will look at in the final part of this tutorial next issue.

Styling the page

With our simple CSS block and DIVs all set up, we will now look at styling the page and some different layout types. This should merely involve trying different CSS definitions, and you will be surprised just how much you can format a page. With our current design in **Fig1**, it would be true to say that our page is looking rather simplistic, and the colours are incredibly bland. We will work on each DIV in turn, and begin with the title.

In our title DIV we really want the title to stand out and look bright. We will first change the background colour to a nice tasteful mid-tone blue:

```
background: #90AEFF;
```

Next we need to make the main title text larger and change its colour. Before we do this however, I think it will look nicer if we change the font. You can do this with the font-family definition:

```
font-family: verdana, trebuchet ms, sans-serif;
```

With this definition we specify the order of font choices for the text. In our example, if **verdana** is on the system, that will be used. If **verdana** is not available, **trebuchet ms** will be used, and as a last resort, **sans-serif** will be used. We can now increase our font size by using the font-size definition:

```
font-size: 32px;
```

This definition states that we want our font to be **32** pixels in size. We can change the colour of the font by using the color definition:

```
color: #EDF9FF;
```

We have set our colour to a lightish blue. To finish off adjusting this font, we will add a little spacing in between each letter. You can do this with the letter-spacing definition:

```
letter-spacing: 5px;
```

With this definition we are putting a **5** pixel padding in between each letter. Speaking of padding, you may have noticed that the text in the title DIV is pushed right up against the side. It would look better if we add some space around the text. This can be done with the padding definition:

```
padding: 10px;
```

With the title DIV now complete, we will move onto the navigation DIV. The first thing we need to change with this DIV is the background colour and padding. I have used:

```
background: #000000;
```

```
padding-top: 3px;
```

```
padding-bottom: 3px;
```

```
padding-left: 10px;
```

As you can see in this example, you can make use of **padding-left**, **padding-bottom** and **padding-top** to specify the padding for each area around the text (there is of course **padding-right** also). Our main navigation links are a horrible blue colour at the moment, so we will now change these link definitions. We will set some link definitions for the entire site:




Fig2 Same page content as in Fig1, this time styled with CSS.

```
a:link
```

```
{
  color: #DEDEDE;
  font-family: helvetica, trebuchet ms, sans-serif;
}
```

```
a:visited
```

```
{
  color: #ABABAB;
}
```

```
a:hover
```

```
{
  color: #EEEEEE;
}
```

With these link definitions we can set the formatting for a normal link (**a:link**), a link that has been visited already (**a:visited**) and when the mouse is hovered over the link (**a:hover**). To ensure that the dashes in between the links are visible, you will need to add a color definition to the nav DIV such as:

```
color: #ffffff;
```

The final area to customise now is the main content area. We will first set the font, padding and background:

```
background: #ffffff;
padding: 10px;
font-family: verdana, trebuchet ms, sans-serif;
```

To make this area stand out a little more, we will add a border around the edge of the DIV. We can set the type of border (**solid**, **dashed** etc), the thickness and the colour. We will use a **dashed** border for the bottom border and a **solid** line for the left and right borders. The definitions for this are:

```
border-left: solid thick black;
border-right: solid thick black;
border-bottom: dashed thick black;
```

You can see our considerably better-looking page in **Fig2** at the top of this column. Our current CSS definitions are:

```
#title
{
  position: relative;
  background: #90AEFF;
  font-family: verdana, trebuchet ms, sans-serif;
  font-size: 32px;
  color: #EDF9FF;
  letter-spacing: 5px;
  padding: 10px;
```

BROWSER DIFFERENCES

Most people who browse the web regularly are aware of some of the small differences between web browsers, but when you start developing web pages, you are able to join the web browser frustration club. The problem is that there are many different browsers, and some of them don't implement CSS fully or correctly. Not only this, but the actual quality of output for the code tends to vary among browsers.

I recommend that you develop your pages for *Mozilla*, as it is a standards-compliant browser. Although it is a good browser to develop for, ensure that you test your code on different browsers. In this series I am designing exclusively for *Mozilla*, so you may see some differences with other browsers, but in the real world, you would need to test your code in all browsers that you can lay your hands on.

```

}

#nav
{
position: relative;
background: #000000;
padding-top: 3px;
padding-bottom: 3px;
padding-left: 10px;
color: #ffffff;
}

#content
{
position: relative;
background: #ffffff;
padding: 10px;
font-family: verdana, trebuchet ms, sans-serif;
border-left: solid thick black;
border-right: solid thick black;
border-bottom: dashed thick black;
}

a:link
{
color: #DEDEDE;
font-family: helvetica, trebuchet ms, sans-serif;
}

a:visited
{
color: #ABABAB;
}

a:hover
{
color: #EEEEEE;
}

```

Changing the layout

At the moment our page is looking good, but we may want to move the different DIVs around a little bit. To do this, we can make use of some additional CSS definitions. To begin with, we will put the navigation area down the left side of the page. Before we do this, we will need to change the contents of the navigation DIV so the options are stacked on each other. We can use an Unordered List to do this. This uses the `` block tag and each list item `` is placed within it:




Fig3 Changing layout with CSS is a straightforward exercise.



Fig4 The navigation box aligned to the right.

```

<ul>
<li><a href="home.html">Home</a></li>
<li><a href="computers.html">Computers</a></li>
<li><a href="services.html">Services</a></li>
<li><a href="contact.html">Contact</a></li>
<li><a href="links.html">Links</a></li>
</ul>

```

To get the navigation area to go to the left side, we need to first set the size of the DIV. We can set this to 200 pixels with:

```
width: 200px;
```

Next, we need to set the float definition to **left**. This definition allows you to specify how the DIV is aligned with respect to the other DIVs. With this definition set to **left**, the content DIV with flow to the right of it:

```
float: left;
```

When you now look at the page, the navigation DIV appears to have disappeared. Rest assured – it is still there, but it is sitting under the content DIV. A simple solution to this problem is to set a **left** margin on the content DIV to the size of the navigation DIV. We set the navigation width to 200 pixels, so we set the **margin-left** of the content DIV to the same:

```
margin-left: 200px;
```

With this set, our layout has been changed, and can be seen in **Fig3**, below left. If you would prefer the navigation on the right side of the page, you can set the float definition to **right** in the navigation DIV, and the **margin-right** definition in the content DIV to **200px**. As you can see, you can really change the layout of the page by just changing key parts of the CSS. This is illustrated in **Fig4** at the top of this column.

Conclusion

In this issue we have started looking at CSS, and how it can be applied to our content to style it. There is no doubt that CSS is a powerful technology, and we have only scratched the surface of what can be done with it. The great thing about CSS is the sheer depth of flexibility that it offers. There are countless definitions and settings that can be applied to your code, and many good-looking designs can be achieved without relying on a single image. Not only does this make your site readable to more people, but it makes it quicker to download and requires lower bandwidth. When we are browsing the web, there are so many sites that we see that rely on images to do what CSS can do without them. [LXF](#)

NEXT MONTH

We will complete this series by looking at some of the advanced features in CSS and how to create some graphically impressive designs and how to centralise our CSS coding.

3D MODELLING

Mesh techniques

PART 10 Jono Bacon continues his *Blender* exploration, this month looking at how to model complex objects.

So far in our *Blender* series, we have looked at a great many techniques for creating 3D scenes, animations and games in *Blender*. This month we are going to look at how we can model complex objects such as animals, humans and other entities. This kind of modelling is of the type that is used in movies such as *Finding Nemo* and *Toy Story*, but of course on a much simpler level for our tutorial! Although creating these complex objects can take time to get right, the actual process of creating the object is relatively simple and requires only a few techniques to create something worthwhile. We will be looking at these techniques and then applying them step by step to create a simple animal character as shown below in **Fig1**.

Skin & mesh

Put your hand out in front of you and look at it. Go on; don't worry if that guy opposite you on the train is going to look at you a bit strangely, just put your hand out, stretch your fingers and look at it.

If you look carefully you will notice that your skin is essentially stretched over the bits inside your hand such as the bones, tendons and muscle. If you were to imagine your hand without the skin on it (not nice, I know), you would have a similar shape (your hand) but would lose the curvature that your skin gives on your

Fig1 A strange fish-like, duck-billed aquanaut that we will create in this *Blender* tutorial.




ILLUSTRATION BY SHIGETO MAEDA - www.blender3d.org




Fig2 Our beginning mesh for creating the aquanaut – a tapering cylinder.

hand. In this sense, the skin on your body is what essentially gives you shape, as it is stretched over your bones and organs.

This concept of skin stretching over something is how *Blender* approaches modelling complex objects. Previously in our modelling exercises, we were creating shapes that were akin to woodworking; we took a shape, resized it, extruded it and moved it. In this issue we will be creating a theoretical mesh over which a notional skin is stretched. With this combination we can not only create natural-looking objects, but ones that can be animated to move in a convincing fashion as well. You will notice a substantial difference in your modelling using these techniques.

Starting modelling

To begin modelling our aquanaut, we will start by creating the mesh that the skin is stretched over. To do this, we will create a basic mesh outline and then we will start to sculpt our animal. First start with a new project, and split the view into a large front view, a small isometric view (you can get this view by pushing *both* mouse buttons together and moving the mouse), and a small side view. Next (in the Front view) add a Surface Curve by opening the toolbox with the **Spacebar** and selecting Add> Surface-Curve. If you zoom into the curve, you will see a small curve added within a number of lines. These lines are used to control the curve, and the curve acts as the mesh in which the skin is stretched over. Now close the curve by pressing the **C** key. You should now have a squashed-looking circle. You can now scale it up a little if needed by pressing the **S** key.

Next, leave edit mode and press **Shift-D** to duplicate your curve. In the side view, scale the size down a little and move it a few blocks away from the first curve. You should now be able to see the smaller curve through the larger one in the front view. With these two curves, we now want to combine them to form our main mesh. To do this we first need to join them together and

then we can add the skin to the mesh. In non-edit mode press **A** so both curves are pink in the side view. Now press **Ctrl-J** and select Join Selected Nurbs from the box that pops up. If you now press **Tab** to enter edit mode, all vertices will be shown (they are currently unselected). Select all these vertices by pressing the **A** key and they will all turn yellow. Now press **Ctrl-F** and you will see the mesh applied to the curves.

In *Blender* there are two important types of lines that are used to model shapes; U and V lines. These special lines refer to the direction in which a shape is subdivided into segments. It is these lines that are used to shape your object, and V lines are the lines down the side of the object with the U lines being the rings. In the Edit buttons now press the Endpoint U button to make the first curve a bit easier to deal with. If you now go to the Edit buttons, you can adjust the resolution of U and V lines by setting the Resol U and adjacent V buttons.

I recommend a setting of 24 for both buttons. You should now have something akin to that of **Fig2** on your screen. **>>**

BIRTH OF A BLENDER COMPETITION

Test your skills against other *LXF* readers

Next month we will be finishing off this series, and to test your knowledge of *Blender*, we will be instigating a competition to see which reader can develop the most impressive piece of *Blender* art from a given task that will be revealed in the next issue. Although we will not reveal the task now, we thought it would be fair to give you all a head-start and fine tune those *Blender* skills in preparation for the competition. It is recommended that you practice all aspects of the series (apart from the game engine and animation articles, this will be a still

image competition) but pay particular attention to lighting, materials and most importantly modelling. The techniques discussed in this very issue will be of particular importance in the exercise set by the competition.

As for prizes, we will have some good *Blender* swag up for grabs for the winner with some additional runner up prizes as well. The winner will also get a high-quality, framed edition of their art, and will be announced and displayed in a future issue of *Linux Format*. Fire up *Blender*, stretch your fingers and get blending!

TUTORIAL Blender





Fig3 The model is rapidly starting to take shape.

Beginning to model our shape

With the mesh now in place, press the **Z** key in the isometric viewport; you will now see a hollow mesh. When we are modelling our shape it is essential that we try to keep our shape symmetrical. The reason for this is to keep the shape simple and easy to subdivide later when we shape our animal. To begin, select the front top-left vertice in edit mode. With the vertice selected press **Shift-R** and you will see a line of the shape highlighted. If you press **Shift-R** a few more times you will see that it will flick between two different lines. These lines are the U and V lines discussed earlier, and **Shift-R** allows you to toggle between the two easily. This means that if you want to edit a particular part of the shape, you select a vertice in that area, press **Shift-R** and model it.

To actually model the shape, we will need to select the relevant portion, subdivide it and model it using the normal tools such as move, scale, rotate, extrude etc. We will first create a mouth by pressing **Shift-R** so the front side of the shape is selected. Now press the **E** key to extrude the shape and then instantly press the **S** key to scale the extrusion smaller. You will

Fig4 The head begins to become recognisable as part of the creature that we saw in Fig1.



see in the isometric view that the side of the shape will close up and you can create a small oval hole. Next press the **G** key to move the oval down a bit so it is in more of a mouth position. If you now move to the side view, you can rotate the currently selected vertices slightly and move them inside the shape a little so the mouth appears recessed. Your model should look somewhat similar to that shown in Fig3, left.

Creating features

We will now make a 'nose' or 'beak' on our shape. To do this we need to select one of the outside vertices on the front of the shape, press **Shift-R** to select that cross-section, then hold shift, select a vertice on the mouth frame and then press **Shift-R** to select the frame. We now have the outside front frame and mouth frame selected so we can create another frame in-between these. To do this, press the **W** key and select Subdivide. You will now see an additional frame appear between those selected. You will see that there are two vertices above the mouth on the new frame. If you select these two vertices, you can press **S** to scale them a little smaller and move them back (using the side view). If you now select the outside frame and the new frame and subdivide, you can extrude a top part of the nose. You can now point the nose up slightly. this will require selecting individual vertices and using **G** to move them.

When you create each new feature, always remember to select the symmetric vertices, and adjust the vertices near the feature to give as much detail as possible. You should subdivide as little as possible, as each subdivision will add more complexity to the shape. The goal is to create a shape that you want while keeping subdivision as low as possible. At this point you can adjust the shape of the mouth and nose using these techniques to get the proportions right.

BLENDER 2.3

New version, updated interface

By the time you read this tutorial, *Blender 2.3* will be available to download. As I write this however, 2.3 is not released, but thanks to the work of blenderhead Ray Wells, I managed to get my hands on a copy of 2.3 just before it was released. As you can see from the screenshot Fig6 above right, there have been quite a few changes made to the release, some of which in the user interface department. There was originally going to be a 2.29 release of *Blender*, but before it was to be released it was decided that development would continue onto the 2.3 series to implement the much-needed user interface improvements. These improvements are considered critical by the *Blender* developers and the first set of improvements is in the 2.3 release.

The main user interface updates have been made to the layout of the buttons and their location. Instead of dominating the user with all the buttons at once, they have now been placed in categories such as Shading for the materials, textures etc. Each viewport is now more clearly defined and the viewports look much clearer. Menus have been tidied up and there is now the addition of a help menu to provide access to documentation from the *Blender* documentation project.

There are many features that have been added in terms of functionality such as 3D font support, *Quicktime* support for textures, theme support, mousewheel support, mirroring, improved vertice handling and the all important undo feature which has been requested by users for many years. It is highly recommended that you give the new *Blender 2.3* a try, but some of the interface descriptions in this series will not apply exactly, as 2.3 has moved many interface elements around to work more efficiently.




Fig6 **BLENDER 2.3**: the mesh gets more complex as more objects are added to the model. Regular readers of this series will note some of the differences evident in the GUI. There's also a Help menu added, along with many more tweaks and improvements.

NEXT MONTH

Next issue is the final part of this series, and we will have a special six-page bumper article to finish off the tutorials properly. We will be looking at some further modelling techniques for creating complex shapes, and look at how to add some more life to your characters.

With our nose pointing upwards, we can now select the top vertices, press **Shift-R** and select the back vertices with **Shift-R**. If you subdivide these frames, you can get another frame that can be moved upwards to form the back of the head shown in **Fig4**, lower left on the opposite page.

To make the shapes I have simply added two mesh UVSpheres and sunk them into the main shape. You can also subdivide the area where the eyes are located and create a ridge for the eyes to be sunken into by scaling the subdivision down and moving it back in the side view. The process should look something like **Fig5** on your screen.


To develop the body, you will need to select the back frame and stretch the body out. If you select this frame and the main head frame, you can then subdivide the body and fatten the beast out a little bit. When the body is fat enough, you can extrude the back frame and rotate it slightly to give a curl in the tail. To finish off the animal, I have added a number of fins to give him the look of a (pretty odd) sea creature. To do this you can just add a cone, press **S** to scale it and press both mouse buttons in top view to flatten the cone to be a simple fin. You can then place the fin on the animal, slightly rotate it in the side view and duplicate three times.

You can of course add any features that you want to the animal. I suggest that you maybe add some ears (well, you can buy 'fish fingers' so why not?), teeth, side-fins, detail to the eyes or other small features. It is the tiny details that are important when modelling creatures, and particular attention should be shown to the eyes. Eyes are possibly the most important feature because the eyes give a creature the life – just take a look at someone with their eyes shut – they just don't look as 'real' without the definition of the eyeballs. This can be obtained to an extent in the materials used when you render the animal, and you could create an eye texture to wrap around the UVSphere.

Summary

In this article, we have taken our first steps at doing some practical modelling of complex shapes. This article has merely provided an introduction to this topic, and we will be looking at some more techniques in the next issue. Although we have covered some simple techniques, it is important to realise how much power these simple techniques now give you. We have essentially covered how to create virtually any shape in *Blender* using these simple techniques. It is this kind of modelling that will help you create impressive and expansive scenes in *Blender*. **LXF**

Fig5 The eyes added to the aquanaut.



TUTORIAL GIMP



PHOTO EDITING

Filter Fantasies: Part 1

Wading through the GIMP menus can be daunting, and finding a filter that fits your needs can be near impossible. But help is on the way, as **Michael J Hammel** retreads some ground to consolidate your skills.

Up till now we've meandered in our tutorials, never really backtracking to the basics. The assumption is that you have found your way to *The GIMP* and have at least one text telling you where everything is (including that ever illusive Canvas Menu).

But perhaps its time to dig a little deeper into the one area most GIMP users would like to know more about: GIMP Filters. In this tutorial (and next month's one as well) we'll be looking briefly at the types of filters available, as well as examine some specific filters and why you might want to use them to solve particular problems or create attractive effects.

LAST CHANCE

Send in your GIMP questions

If you've been following the LXF GIMP tutorial series, and there's something that we've covered that you haven't understood or would like us to expand upon, please send an email outlining your concerns to our usual linuxformat@futurenet.co.uk

with 'GIMP Question' as the subject line. Or maybe you feel that there's an important topic that we've left out? Questions both large and small will be dealt with in a special Q&A episode in a few issues' time, so get in touch before it's too late...

ARTISTIC MYOPIA: DEPTH OF FIELD

Depth of field is the distance between the nearest and farthest points that appear to be in focus in a photograph. If an object is 10 feet away and the depth of field is between 8 and 12 feet, then the object would be in focus. Any objects outside that range would be out of focus.

If a photograph was taken using a large depth of field, then objects in both the background and foreground may be in focus, possibly making it difficult to distinguish the true subject of the photo. To draw attention to that subject we can use basic *GIMP* filters to fake a smaller depth of field. We start by scanning a photograph using *XSane* or *Quitelnsane*, open source applications that make use of the SANE drivers for scanners.

Original vs sharpened

1 Most scanned photographs tend to be a little out of focus to begin with, so we apply the **Filters>Enhance>Sharpen** filter to clean up the scan. We could perform this step after changing the depth of field, but doing so would also sharpen the areas just taken out of focus. So sharpen first!



Quick Mask and feather

2 The next step is to make a loose selection around the subject of interest. In this case the dog in front is to be our focal point. We use the Quick Mask tool (lower left of the Canvas window) to paint over the subject – anywhere the image shows through the red veil will be the area to be selected. Don't worry about being exact here. After we're done painting, we convert the Quick Mask to a selection (the button next to the Quick Mask button) and then invert it (**Selections>Invert**). The selection is then feathered a modest amount (10 pixels in this example). The feathering here is why we did not need to be so exact while making our Quick Mask selection.



Gaussian blur = final image

3 With the background selected we choose a **Gaussian Blur (Filters>Blur>Gaussian Blur IIR)**. Note that the IIR blur is better-suited to photographs than the RLE version (which is better suited for computer-generated and text images). The blur amount is subjective – choose the amount which takes the non-focal subject area out of focus. Any more than that becomes a lesson in artistic license.

RED IS GREEN, GREEN IS CLEAR: COLOUR MAPPING

Technically, colour mapping is the process of reassigning one colour value to another. Each colour channel (Red, Green or Blue) has 256 distinct levels in *The GIMP*. $256 = 8 \text{ bits} = 1 \text{ byte}$ per colour channel, which is convenient for the *GIMP* developers and extremely efficient for computers to handle. This is why *GIMP* is known as an 8-bit graphics app – 16 bits would give greater colour variation, though the average desktop user isn't likely to notice the difference.

Mapping colours doesn't have to be a one-to-one process – you can map 1000 colours down to 256. This is exactly what happens when you convert a JPEG image to a GIF, the former offering thousands of colours while the latter being limited to 256. In *GIMP* you can convert a Canvas from an RGB image (256 colours per channel) to an Indexed image, which has only 256 colours total, prior to saving it as a GIF image. This would be colour mapping in its simplest form. Colour mapping can also be artistic. Most of the options available under the Filters>Colours menu offer some form of colour mapping that can take ordinary images and turn them into something far more imaginative. There are many other ways to manipulate colour in *The GIMP* besides those filters found under Filters>Colour. Try working with the options under Image>Colours, especially the Curves, Hue-Saturation and Filter Pack options. These tools are designed more for colour correction than for artistic use, but with *GIMP* you're only limited by your own creativity...



Alien Map 2

2 While Gradient Map works on the brightness of an image, the Alien Map plug-ins (Alien Map and Alien Map 2) work on the RGB and HSV colour spaces. There are many colour spaces, such as CMYK and YUV, but *The GIMP* does not yet support them all. The RGB colour space is the one you work with in the *GIMP* most often – this is simply working with the Red, Green and Blue channels. HSV is a little different. The image has Hue, Saturation and Value levels than can be mapped directly to RGB. Applying similar methods to either colour space will produce different results. In this example, we've used Alien Map 2 with the same settings but applied first to the RGB colour space of the image and then to the HSV colour space. While the latter may look far more artistic, this is only because of the values chosen for the Alien Map 2 filter. Different values could very well produce better results for the RGB colour space.



Gradient map samples


1 We start with the image we finished with in our depth of field exercise. Duplicate it. We want to apply two variations of a Gradient Map (Filters>Colours>Map>Gradient Map) to it. This filter will map the colours in the image (actually, their brightness) to the colours in the gradient chosen. In our first example we chose the 'German Flag Smooth' gradient. The second example we applied the 'Golden' gradient. You can see the great difference in the result when gradient to be used is chosen carefully. Try applying a gradient with lots of transparency, such as Burning Transparency.

ADDING NOISE TO A CLEAR PICTURE

Noise is a random change in an image. Pixels are chosen and changed randomly by a noise filter, altering colour or brightness or some other aspect of the pixel. There are various reasons to add noise to an image. One of the most common is to remove banding (visible lines of solid colour) in gradients when an image is changed from RGB to Indexed mode. Noise can be used to blend image elements in a photograph such as melding the foreground into the background – opposite of the depth of field, you might say. Noise is also used to simulate movement in GIF animations. We'll look at two filters from the Filters>Noise menu in these next two examples.

Scatter HSV

1 The Scatter HSV filter is just one noise option available to *GIMP* users. Changes in Hue will vary the amount of colour change for the noise. Changes in Saturation will change the intensity of those changes in colour. Changes in Value will alter the brightness of the noise. The example here shows the Scatter HSV dialog and its settings along with the results of applying these values to add noise to the picture.





Spread

2 Another noise option is Spread. The effect produced by this filter is something between a oil painting done by dabbing the brush tip on the canvas repeatedly and viewing an image through translucent (as opposed to transparent) glass. In the example here we see the results of applying a Spread of 30 in both the horizontal and vertical directions to the original image along with a close up to show the detail produced by the filter.





BENDING LIGHT: GLASS AND LIGHT EFFECTS

All photographs are frozen moments of reflected light. *The GIMP* can let you change the perceived direction of those reflections or add your own. The filters found under the Filters>Glass Effects menu change an image to make it appear the image was taken through special lenses or through semi-translucent glass windows. The filters under Filters>Light Effects, on the other hand, provide a means to generate new types of lighting, from point and spot lights to lens flares.

Glass Tile

1 The Glass Tile filter generates patterns similar to what you see with the Spread filter by generating images that appear distorted from diffracted light patterns. Light is diffracted when it hits an edge, so the patterns produced with the Glass Tile filter amount to viewing the image through a wall made of glass tiles of varying sizes. Such tiles are often used in windows for bathrooms to obscure the view inside from the outside, or in open-plan areas as a focal point or room divider. This example shows the Glass Tile filter dialog along with the results after applying it to the original image.



GFlare

2 A lens flare is what happens in photography when a light source (the sun or a light bulb, for example) reflects on a camera lens and causes ghost images in the form of spikes of light radiating from the light source. These spikes can be solid or come in the form of multiple circles of coloured light, depending on the complexity of the camera lens and even on dust particles on the lens. Most lenses and lens filters are designed to try to remove this effect, even though many photographers and software developers like to put them back in. Because a lens flare is produced on the lens it will always appear in front of the objects in an photograph. The GFlare filter can produce a wide variety of lens flares without regard for the existing image content. Note that you should generate these flares on a separate layer and keep this layer on top of your photos for the most realistic look. Lens flares with hot spots (small circles near but off-centre from the centre of the light source) are caused by dust and can be reproduced with multiple layers of GFlare renderings. [LXF](#)

PHP VERSION 5

Practical PHP

PHP 5 heralds a new era for PHP programmers, or so the developers would have us believe. Paul Hudson separates the outstanding features from the outrageous hype and jumps on the bandwagon...

Happy New Year, PHPers – it's January at last, and with the new year we also get a major new release of our favourite programming language. Having been in development for much of 2002 and all of 2003, PHP 5 brings with it a healthy swathe of new features, bug fixes, enhancements, and optimisations that, if carefully harnessed, can and will improve your scripts. The purpose of this article is to kick off a tour of PHP 5, examining what's changed and why, as well as providing a comprehensive introduction to its biggest new features. Before we get started, though, I want to supplement *LXF46's* exclusive from PHP developer Zeev Suraski by briefly outlining exactly how we got to the current situation...

Personal home pages

PHP 3 was the first version to achieve widespread use on the Net, largely because it finally had a reliable script parser – before that, any errors reported in your scripts were as likely to be your fault as it was the developers'! Most of the advantages to PHP 3 are still the key advantages of PHP 5 today: quick and easy script programming, easy database access, and a wide range of built-in functions to let you focus on the hard stuff yourself.

The big changes in PHP 4's lifespan were based around performance and security – the execution paradigm for scripts

was changed from “execute each line as it's read” to “parse first, execute later”, which is a great deal faster. On the security front, the PHP 4 series also saw `register_globals` being disabled by default, which, while causing a lot of anguish from existing programmers, actually serves to make everyone's scripts safer in the long run.

PHP 4 also managed to add a variety of new functions and thousands (literally) of bug fixes, but the biggest difference felt was in speed and security.

What's new?

With PHP 5, the focus is on features, features, features, and quite rightly – we're not likely to see as big a change as PHP 3 to PHP 4 ever again, which has given the developers more time to focus on all the requests they've received for additions to the language. Syntactically speaking, PHP 5 is largely similar to PHP 4.3.4 and you should be able to import your scripts across wholesale.

If anything, the change that is most likely to catch you out is the big reworking of object orientation. Since this feature was introduced in PHP 3, it was really just a hack – you could define classes and instantiate objects, but many features of the OOP paradigm were missing, such as access control, destructors, abstract and final classes, and static class variables.

Of course, there are a variety of other features in PHP 5, but this issue we're covering the OOP changes alone, because, as I mentioned, they are most likely to break backwards compatibility.

In subsequent issues we'll be looking at the new XML module, *SimpleXML*, and also the new flat-file database back-end, *SQLite*.

Handling objects

The biggest difference in the PHP 5 object model is the way in which it handles object references. In PHP 4, in order to keep objects working like other variables, when you assigned one variable to another you performed a copy. That is:

```
$foo = "bar"
$wombat = $foo;
$foo = "baz";
```

In this situation, **\$wombat** would still be set to **bar** even after **\$foo** has been changed to **baz** because the variable data was copied in its entirety and not as a reference. In PHP 4, copying an object did just that – you had a unique and individual instance of that object that you could manipulate independent of the original. However, this led to quite a lot of confusion from most developers, and also left many scripts running a great deal slower than they needed to.

From PHP 5, objects are now referenced via a unique handle that serves as an identifier. An object variable, instead of holding the actual object, now holds the handle number of that object. When the object is used, the handle is resolved to the actual object, but when the object is copied it is the handle that is copied not the object. In PHP 4, this is the equivalent of always copying objects by reference, and means that you need to go to extra lengths if you want to actually copy the object data. More on that later – first let's look how the new reference-by-default action helps.

```
$foo = new myobject;
dobar($foo);
```

In that code, we pass the handle to **\$foo** into the **dobar()** function, which means that any changes made to **\$foo** inside

dobar() are made to the actual object itself. This solves a regular problem encountered by novice programmers whereby copies are passed in and copies are returned by default – a problem that usually confuses a great deal! By referencing objects by default, PHP eliminates a great deal of object copying – traditionally a very slow task – which will therefore help speed up the majority of scripts once any code-breaking upgrade problems are fixed.

If you explicitly want to copy an object rather than reference it, there's a magic function just for you. Magic functions, if you've yet to meet them, are PHP functions built into each object by default that start with two underscores. To copy an object, call the **__clone()** magic function, which returns a copy of the object. Here's an example:

```
$foo = new myobject;
$bar = $foo->__clone();
```

That will create a new instance of the object and return its handle. If you want to update your PHP 4 scripts to work smoothly in PHP 5, you'd use code like this:

```
$foo = new myobject;
dobar($foo->__clone());
```

The **__clone()** function is just the beginning, though – there are quite a few other magic functions introduced with PHP 5 that make things all the more interesting...

Ha ha this-a-way

Magic functions are PHP's way of providing some base functionality to all objects, and might seem a little confusing and unusual at first until you think that this sort of functionality has been around for quite some time in other languages – compare the all-mighty **TObject** in *Kylix*, for example, which includes a selection of basic functions by default.

There are six new magic functions introduced in PHP 5 and they serve a variety of purposes. I've ordered them here easiest first to get you started quickly...

First off the bat there's **__toString()**, a peculiar function that, if you define it, is automatically called by PHP if you attempt to use your object as if it were a string. If you're questioning why **__toString()** is important, there are a variety of possible uses – perhaps the most interesting is to make it print out a selection of information about the object, sort of like a filtered **var_dump()**.

A magic function for classes that is global as opposed to being attached to a class is **__autoload()**, which is another unusual but helpful function that, if you have defined it, is called whenever you try to instantiate an object of a class that doesn't exist. For example:

```
<?php
function __autoload($class) {
    echo "$class: No such class defined!\n";
}

class Dog() {
    function whatever() {
    }
}

$mrjingles = new mouse();
?>
```

Here we create a new mouse object despite only having a **Dog** class defined. The end result is that **__autoload()** is called, with the name of the non-existent class passed in as a parameter, and



OOP: WHY BOTHER?

Object Orientation *not* slow!

Despite the OOP paradigm having been commonly used since *Smalltalk*, many people still shun objects and consider OOP to be slow and unwieldy. It's true that, back in the early days, object orientation was pretty poor in C++ because compilers did a hashed job, but this has improved a long way. In PHP 3 and 4, OOP was also a pretty hashed job, but again it has been improved massively in PHP 5.

OOP is a great way to give a limited amount of intelligence to your code whilst also forming coding contracts with yourself and other programmers. It makes a lot more sense to say **Player->Draw()** and have the **Player** object know how to handle its own drawing than have to have your main **draw()** function know how to draw the **Player** object, simply because it means that if you ever find your player needs to be have some special drawing code that you hadn't anticipated earlier, you only need amend the **Player->Draw()** function – all the other code doesn't care what **Player->Draw()** does because it's all handled internally. Using this thinking allows you to pick up the **Player** class and drop it into other projects easily, which promotes code re-use.

The code contracts idea is also key, because OOP allows you to define what variables and functions can be accessed and where, which means that if you (or someone else) tries to do something that you previously decided shouldn't be allowable, they will be stopped dead in their tracks.

TUTORIAL PHP

◀◀ the code outputs a warning. While this is all well and good, what makes it *very cool* is the fact that, after calling `__autoload()`, PHP will attempt to instantiate the class again. In practice this means that you can attempt to create a class, and, using `__autoload()`, automatically **include()** the script containing the class definition if it's not included already, like this:

```
function __autoload($someclass) {
    include "$someclass.php";
}
```

Using the code above, attempting to create an object of class **mouse** will fail, the PHP will attempt to include `mouse.php`, and finally attempt to instantiate the **mouse** object again.

Build it up

Handling the creation and deletion of objects in PHP has always been a hassle, with various hacks being using by PHP 4 programmers to attempt to simulate a destructor. In PHP 5, this is now fixed with a unified constructor/destructor system. For the uninitiated, a constructor function is called when an object is

created, and a destructor when deleted, and are helpful to initialise and clean up your objects.

Construction and destruction are handled using the `__construct()` and `__destruct()` magic functions respectively, which are called automatically as you work with the object. Take a look at the following script, for example:

```
<?php
class mouse {
    function __construct() {
        echo "This is a mouse\n";
    }

    function __destruct() {
        echo "This is an ex-mouse\n";
    }
}

echo "Stage 1\n";
$mightymouse = new mouse();
```

ACCESS CONTROL

Private, public and protected modifiers

So far we've covered about half of the new OOP features in PHP 5, and hopefully you can already appreciate that it's a worthy upgrade. One of the key advantages to OOP is that it lets you form code contracts, that is, it lets you specify precisely which parts of your scripts may access what variables inside your objects, which acts to re-enforce your design specification at the code level.

There are three types of access control: private, public, and protected, which respectively make variables and functions available only to the class, available to everyone, and available only to the class and classes inheriting from it.

The public modifier was implicitly used by PHP 4, which meant that all properties and functions defined in a class could be used from anywhere in your script. In PHP 5, this would be written like this:

```
<?php
class sheep {
    public function __construct() {
        echo "Baa!\n";
    }

    public function __destruct() {
        echo "Aab!\n";
    }

    public function do_cheese() {
        echo "Baa baa black sheep\n";
        echo "Have you any wool?\n";
    }
}

$dolly = new sheep();
$dolly->do_cheese();
?>
```

As seen in the other scripts shown so far, the public keyword isn't required, because PHP 5 continues to implicitly use public access if you don't specify otherwise. While

this does mean that scripts are compatible with both PHP 4 and PHP 5 simultaneously, it does leave that little bit extra to human error – I'd recommend you specify public explicitly, so as to make your intentions clear.

Now, try changing the `do_cheese()` function to this:

```
private function do_cheese() {
    echo "Baa baa black sheep\n";
    echo "Have you any wool?\n";
}
```

When you run the script this time you should get this error back:

```
Fatal error: Call to private method
sheep::do_cheese() from context "
```

Now that the function is private, we can't call it from elsewhere in the script, because private functions are only supposed to be called from inside the object itself. Similarly, we can declare variables as private, like this:

```
<?php
class koala {
    private $Name;

    public function Name() {
        return $this->Name;
    }

    public function setName($NewName) {
        $this->Name = $NewName;
        return true;
    }
}

$fuzzy = new koala;
$fuzzy->setName("Fuzzy");
echo "{$fuzzy->Name()}\n";
?>
```

Here we have a classic setup in the OOP paradigm – a private variable with public accessor functions that are used to set and get the variable.

The difference between private and protected variables is important to understand, as both have their uses when it comes to contract programming. A privately declared variable is accessible only to the class in which it is declared, which is very useful if you only want the functions of that object to have access to those variables – this stops inherited classes playing around with the parent class. On the other hand, protected variables are available to use from the child classes, which allows much more flexibility at the possible expense of code reliability.

At the time of writing, PHP has a bit of an irregularity in its OOP system. Take a look at this script:

```
<?php
class dog {
    private $Name;
}

class labrador extends dog { }

$dozer = new labrador;
$dozer->Name = "Dozer";
print_r($dozer);
?>
```

That will output this:

```
labrador Object
(
    [Name:private] =>
    [Name] => Dozer
)
```

General consensus is that this is part of the plan because PHP does generally let you define variables on the fly, but it does make things confusing – PHP sets a variable called `$Name` to the value you pass, and many might think that this is it overriding the private declaration. In actual fact it ends up with two `Name` variables in the same object, which is almost certainly not a desired outcome. Be careful!

```
echo "Stage 2\n";
unset($mightymouse);
echo "Stage 3\n";
$dangermouse = new mouse();
echo "Stage 4\n";
?>
```

(Note: No mice were hurt in the production of this tutorial!)

That code will output the following text:

```
Stage 1
This is a mouse
Stage 2
This is an ex-mouse
Stage 3
This is a mouse
Stage 4
This is an ex-mouse
```

There are two things to note there. Firstly, calling **unset()** on an object calls its destructor implicitly as the object is deleted, as you can see at **Stage 2**. Secondly, as PHP performs garbage collection at the end of the script, it will automatically destroy the **\$dangermouse** object and call its destructor in the process, hence the ex-mouse notice at **Stage 4**.

Chain constructions

Using **__construct()** on a simple object is fairly straightforward, as you've seen. However, when you work with classes that inherit from other classes, things get a little more complicated. Consider the following class hierarchy:

Animal



Dog



Labrador

The grandfather class, **animal**, will have various properties that are generic to all classes beneath itself, such as age, height, weight, *etc.* The parent class, **dog**, will have properties generic to all breeds of dog, such as `number_of_lampposts_sniffed`, `tail_wag_level`, *etc.*

Finally, the child class, **labrador**, will have variables unique to the labrador breed such as `number_of_andrex_adverts_starred_in`, *etc.* Yes, this is a little contrived, but you should see the point! For a labrador dog (we'll call him **Oscar**) to exist fully, he needs all the variables from the child class, the parent class, and the grandparent class to exist and be initialised properly.

So, how do we go about doing this? Take a look at this chunk of code, which is the PHP representation of above class hierarchy:

```
<?php
class animal {
    function __construct() {
        echo "Animal created\n";
    }

    function __destruct() {
        echo "Animal deleted\n";
    }
}

class dog extends animal {
    function __construct() {
        echo "Dog created\n";
    }
}
```

```
function __destruct() {
    echo "Dog deleted\n";
}

class labrador extends dog {
    function __construct() {
        echo "Labrador created\n";
    }

    function __destruct() {
        echo "Labrador deleted\n";
    }
}

$oscar = new labrador;
?>
```

Save that script and run it using the CLI SAPI – you should see messages about a labrador being created and deleted. However, what you won't see is any message about a dog or an animal being created and deleted, despite Oscar clearly being both. The reason for this is because PHP will call one and only one constructor, which is **labrador::__construct()** – it won't call **dog::__construct()** or **animal::__construct()**, which means if you have anything more substantial in those two functions that just echoing out a minor notice, your code will break.

How to get around this? Well, inside the object you have direct access to the parent class through the object **parent**, which means you can statically call the **__construct()** function of the parent class from the constructor of the child class. Sound confusing? It's not, take a look at this:

```
class labrador {
    function __construct() {
        parent::__construct();
        echo "Labrador created\n";
    }

    function __destruct() {
        echo "Labrador deleted\n";
        parent::__destruct();
    }
}
```

The first action that **labrador::__construct()** does is to call the **__construct()** function of its parent class, **dog**. This runs **dog::__construct()** for the **labrador** object, setting up an variables defined at that level (`number_of_lampposts_sniffed`, *etc.*), then goes ahead and runs **labrador**-specific stuff. It's crucial that you call the **parent** constructor before going ahead and running **child** tasks because it's often the case that the **child** class will use properties of the **parent**, and trying to use things before they have been setup is, of course, bad.

There are two further things to note in this area. Firstly, you'll need to add **parent::__destruct()** and **parent::__construct()** to the dog class also, so that it will in turn call **animal::__construct()** and **animal::__destruct()** – this completes the chain. Secondly, note how **parent::__destruct()** comes *after* the **child** code. This is for similar reasons to why **parent::__construct()** is before the **child** code – if you destroy variables created by the **parent** before you're certain you're finished using them, there's a chance you might run into trouble. Good programming practice is to always put **parent::__destruct()** after the **child**-specific code. [LXF](#)

NEXT MONTH

There are three magic functions we've not yet looked at, which are **__get()**, **__set()**, and **__call()**, as well as static class variables and abstract/final classes – we'll be looking at these next month, as well as the *SimpleXML* extension.

EVERYONE SHOULD HAVE A LINUX SERVER

SERVER SCHOOL:

Serving up ftp

PART 5 Nick Veitch unravels the mysteries of the File Transfer Protocol.

FTP stands for File Transfer Protocol, and is one of the backbones of the Internet. Long before HTTP came along with its fancy graphical interface, FTP was the way people exchanged information from one server to another. It is still very popular today, mainly because, as opposed to HTTP-based file downloads, it is much simpler, both from the client and the server side. Many mirror sites use it simply because the overhead of running webservers simply to server up pages listing filenames isn't worth it.

Properly set up and configured, FTP can also be more secure and reliable than a web-based file service, and of course it can

combine with other methods of remote access, such as *NFS* or *Samba*, to provide an intra- or inter- networked solution.

Security

One of the most important things to bear in mind with all FTP servers is that of security. Because the daemons need to run as root, there is a certain risk that any FTP server could compromise your system. This is where, for simple downloading of files, serving via HTTP can be an advantage if you are already running a webserver – you only have one server to worry about.

When a remote user contacts the FTP server, they need to log on to proceed any further. The username and password are generally transferred via clear text across the network – another security risk. The FTP server then checks these against the users registered on the local system and changes to an appropriate directory. This means that users on the local system (including root) can login and transfer files to and from their own home directories – another potential security hole unless properly managed.

Furthermore, the FTP protocol allows the client to execute commands on the server. This leaves the door open to exploit known weaknesses in other software on the system. Fortunately there are plenty of ways to tighten up security, by first creating a secure configuration, and secondly by making sure the rest of the system is geared for security.

Anonymous access

One of the first questions you need an answer to is whether anonymous access should be allowed. Obviously for a public-

ftp is much more secure and includes many similar commands to ftp

```

ftp> connect
ftp> login
ftp> ls
ftp> cd
ftp> pwd
ftp> mget
ftp> mput
ftp> bye
  
```

facing FTP site, this is required, though this does make securing it a little harder, though not impossible.

The first step is to create a user account for the FTP daemon. This is usually done when the software is installed, but it depends on the distribution or the packages you use to install. If the user 'ftp' does not exist, create one. As a further precaution, you should set the default shell to something that can't actually be used, like the command **false**.

Configure anonymous access on WU-ftp

The first thing to do is add or edit the user ftp on your system. This user account is usually created when the ftp server software is installed, but may not be. You can create the user any way you like, but you will need to edit the /etc/passwd file on the server to look something like this:

```
ftp:401:401:Anonymous FTP:/home/ftp:/bin/false
```

Note that the default shell has been changed as an extra security precaution – it now points at an executable which won't allow the user to run any commands.

You should also create a group for the ftp user:

```
groupadd -r ftpgroup
```

Now we need somewhere for the ftp users to access. Create the following directories:

```
/home/ftp
```

```
/home/ftp/bin
```

```
/home/ftp/lib
```

```
/home/ftp/etc
```

```
/home/ftp/pub
```

```
/home/ftp/incoming
```

The directories are necessary because we are going to lock the ftp user into this specific directory. All the files they need to access, and the commands they need to use will be contained in here.

You can now copy your passwd and group files from /etc/ into /home/ftp/etc/ and edit them. These are only used for the users and groups that will need to access files via ftp, so you can remove all the other entries. Change all the password fields to a *. The resulting file should look something like this:

```
root:0:0:Admin user::
```

```
ftp:401:401: Anonymous FTP::
```

There are a few more things to do to get the system working.

Copy the **ls** command from /bin to our new ftp area.

```
cp /bin/ls /home/ftp/bin
```

For the sake of pleasantries, you should also create the file 'welcome.msg' in the /home/ftp directory. This is merely a text file containing a message which is broadcast to anyone connecting to the server using a text-based client.


Now we need to set the permissions of the whole file structure. This is important to allow users access only to what they need, and also to prevent them escaping into the system as a whole. Table 1 at the top of the following page shows the file structure, ownership and flags of all the files in the directory. The incoming directory is kept write only as a further precaution to prevent files being seen and taken advantage of by the unscrupulous.

Finally, you will need to make some changes to the /etc/ftppass file to allow uploads to the directory.

```
class all real,guest,anonymous *
```

```
email root@localhost
```

```
loginfails 5
```



The CERT site is worth visiting often. Any new vulnerabilities will usually be flagged up here first.

```
readme README* login
readme README* cwd=*
```

```
message /welcome.msg login
message .message cwd=*
```

```
compress yes all
```

```
tar yes all
```

```
chmod no guest,anonymous
```

```
delete no guest,anonymous
```

```
overwrite no guest,anonymous
```

```
rename no guest,anonymous
```

```
log commands anonymous,guest
```

```
log transfers anonymous,guest,real inbound,outbound
```

```
log security anonymous,guest
```

```
log syslog+xferlog
```

```
shutdown /etc/shutmsg
```

```
passwd-check none warn
```

```
anonymous-root /home/ftp all
```

```
autogroup ftpgroup
```

```
upload /home/ftp * no
```

```
upload /home/ftp/home/ftp/incoming/* yes root root 0664
```

```
nodirs
```

```
private no
```



Many thanks to Rackspace Managed Hosting who have donated hardware to help us complete this series. For more info about Rackspace, see page 94.

◀ TABLE 1: Directory Structure

file	owner	group	permissions
/home/ftp	root	ftpgroup	555
/home/ftp/bin/	root	root	111
/home/ftp/bin/ls	root	ftpgroup	111
/home/ftp/etc/	root	ftpgroup	111
/home/ftp/etc/group	root	ftpgroup	444
/home/ftp/etc/passwd	root	ftpgroup	444
/home/ftp/incoming	root	ftpgroup	1733
/home/ftp/pub/	root	ftpgroup	2555
/home/ftp/welcome.msgroot	ftpgroup	111	

ProFTP

ProFTP is a little more recent than many other ftp servers, and first emerged in the late 1990s. Although it performs the same job as other servers, it takes a very different approach to features and configuration, more closely following the modular design of *Apache* than the monolithic structure of many other FTP servers. RPM and .deb packages are available from the main website, and this is by far the easiest way to install the software.

It is also pretty easy to set up proFTP to accept anonymous connections and allow uploads. Here is an example configuration

```
# This sample configuration file illustrates configuring two
anonymous directories, and a guest (same thing as anonymous
but requires a valid password to login)
```


```
ServerName "ProFTPD Anonymous Server"
ServerType standalone
```

```
# Port 21 is the standard FTP port.
Port 21
```

```
# If you don't want normal users logging in at all, uncomment
this next section
#<Limit LOGIN>
# DenyAll
#</Limit>
```

```
# Set the user and group that the server normally runs at.
User nobody
Group nogroup
```

```
# To prevent DoS attacks, set the maximum number of child
processes to 30. If you need to allow more than 30
concurrent connections at once, simply increase this value.
```



ProFTPD adopts a modular approach, much like *apache*, but is still completely compatible with the FTP standards.

```
# Note that this ONLY works in standalone mode, in inetd
mode you should use an inetd server that allows you to limit
maximum number of processes per service (such as xinetd)
```

```
MaxInstances 30
```

```
# Set the maximum number of seconds a data connection is
allowed to "stall" before being aborted.
```

```
TimeoutStalled 300
```

```
# We want 'welcome.msg' displayed at login, and '.message'
displayed in each newly chdir'd directory.
```

```
DisplayLogin welcome.msg
```

```
DisplayFirstChdir .message
```

```
# Our "basic" anonymous configuration, including a single
upload directory ("uploads")
```

```
<Anonymous -ftp>
```

```
# Allow logins if they are disabled above.
```

```
<Limit LOGIN>
```

```
AllowAll
```

```
</Limit>
```

```
# Maximum clients with message
```

```
MaxClients 5 "Sorry, max %m users — try again later"
```

```
User ftp
```

```
Group ftpgroup
```

```
# We want clients to be able to login with "anonymous" as
well as "ftp"
```

```
UserAlias anonymous ftp
```

```
# Limit WRITE everywhere in the anonymous chroot
```

```
<Limit WRITE>
```

```
DenyAll
```

```
</Limit>
```

```
# An upload directory that allows storing files but not
retrieving or creating directories.
```

```
<Directory incoming/>
```

```
<Limit READ>
```

```
DenyAll
```

```
</Limit>
```

```
<Limit STOR>
```

```
AllowAll
```

```
</Limit>
```

```
</Directory>
```

```
</Anonymous>
```

```
# A 2nd anonymous ftp section. Users can login as "private". Here
we hide files owned by root from being manipulated in any way.
```

```
<Anonymous /usr/local/private>
```

```
User bobf
```

```
Group users
```

```
UserAlias private bobf
```

```
UserAlias engineering bobf
```

```
# Deny access from *.evil.net and *.otherevil.net, but allow all
others.
```

```

<Limit LOGIN>
Order deny,allow
Deny from .evil.net, .otherevil.net
Allow from all
</Limit>

# We want all uploaded files to be owned by 'engdept' group
and group writable.
GroupOwner engdept
Umask 006

# Hide all files owned by user 'root'
HideUser root

<Limit WRITE>
DenyAll
</Limit>

# Disallow clients from any access to hidden files.
<Limit READ DIRS>
IgnoreHidden on
</Limit>

# Permit uploading and creation of new directories in
submissions/public

<Directory submissions/public>
<Limit READ>
DenyAll
IgnoreHidden on
</Limit>

<Limit STOR MKD RMD XMKD XRMD>
AllowAll
IgnoreHidden on
</Limit>
</Directory>
</Anonymous>


# The last anonymous example creates a "guest" account,
which clients can authenticate to only if they know the user's
password.

<Anonymous ~guest>
User guest
Group nobody
AnonRequirePassword on

<Limit LOGIN>
AllowAll
</Limit>

# Deny write access from all except trusted hosts, if required.
<Limit WRITE>
Order allow, deny
Allow from 10.0.0.
Deny from all
</Limit>
</Anonymous>

```



Real security – *sftp*

As we discussed earlier, the main problem with ftp is that the passwords are passed in a clear text format over the network, and are therefore susceptible to being captured and compromised, even if only used on a local network.

This is also the problem with the telnet remote shell, which is why it has fallen into disuse, and been largely replaced by *SSH*, the secure remote shell. Similarly, *SSH* also provides FTP capability, via the provided *sftp* server and client. Essentially, this will allow you to do any of the things you can currently do with normal ftp, but using a secure connection, so you can rest easy that user accounts aren't being snooped upon. The *sftp* client and server is provided by the *openssh* package, which you can retrieve for yourself from your distro vendor or from the main site at www.openssh.com

The only drawback is that there is no flexibility to this system. Obviously you won't be able to safely allow anonymous access if that's what you are after. It also isn't possible to herd all *sftp* connections into the same filesystem. The range of commands isn't quite as extensive, but if you have ssh installed anyway, you can log in to a shell account and do what you need anyway.

sftp is also supported by several Linux GUI clients, so it should even be fairly transparent for end-users too. Ultimately though, *SSH* and *sftp* have also had some vulnerabilities over the years. While it is certainly more secure than a standard ftp connection, you shouldn't kid yourself that you are immune.

Precautions

No matter how secure, there is always the potential that someone will be able to exploit a hole in your configuration or a vulnerability in the software you are using. With something like FTP, the results can be devastating. There are a few things worth considering:

- 1 Locate the ftp filesystem on a separate device and impose quotas. Some attacks on ftp sites have stemmed from the intentional or unintentional capabilities of uploading files. You don't want your main filesystem bugged up with this nonsense if you can help it.
- 2 Run some sort of tripwire software, just to check that your system files are not getting tampered with.
- 3 Turn on all the logs, and review them.
- 4 Consider running your ftp as a properly **chrooted** process. This may require changing your *inetd* daemon.
- 5 Alternatively, and perhaps more extremely, consider running your ftp server under an instance of UserModeLinux. [LXF](#)

Many GUI ftp clients can make use of the *sftp* protocols for file transfers.

NEXT MONTH


Here's a list of the subjects we will be tackling over the coming months in the *LXF Server School* tutorials:

- PART 6** Mailservers
 - PART 7** Databases
 - PART 8** Intrusion detection
 - PART 9** LDAP
 - PART 10** Your Questions
- If there's anything you'd like to see in particular in Part 10, drop us a line to the usual address with "Server School Suggestion" as the subject-line.

KNOW YOUR NETWORK

TCP/IP network infrastructure with Linux

PART 1 When it comes to building computer networks, Linux has just about everything you'll need. Chris Brown begins a three-part look at the power behind the Internet.



Whether you want routers or name servers, firewalls or filesharing, web servers or virtual private networks, Linux has the software to do it. It is not surprising that the greatest penetration of Linux into the commercial world has been as a provider of network services.

In this three-part series, we're going to take a look at the roles Linux can play in providing not only application services, but also the underlying infrastructure of TCP/IP-based networks. In this first part, we'll take a look behind the scenes at how the Internet actually works and what holds it together. Later in the series we'll look at the specific features and applications provided by Linux which satisfy these needs. Linux really is a complete networking solution in a box.

The Internet exposed

Do you have any idea what actually happens when you type in a web address such as <http://www.linuxformat.co.uk/somefile> to your web browser? Well, do you? Tracing through the sequence of events behind the scenes can reveal a lot about how the Internet and the world wide web works.

First, the browser picks apart the URL into pieces as shown in **Fig1**. The first part, **http**: tells the browser that it needs to connect to a web server and communicate using a protocol called the hypertext transfer protocol (http).

The second piece: **www.linuxformat.co.uk** is the name of the computer where the web server is running. If you want to impress your friends you can call it the Fully Qualified Domain Name. Finally the third piece **somefile** is the name of the file you're asking the web server to send to you. In the simplest case, the file will just contain static HTML code which your browser will render to the screen.

Wait, what's a protocol?

We just described http as a 'protocol', and as the IT world appears to love protocols with a passion, we're going to be running into a good few of them in this series, so it's worth taking a minute to explain what they are.

In fact, we're all used to following all sorts of protocols in our daily lives, we just don't call them that. Suppose you want to book theatre tickets. You call up the booking agency, tell them which show you want to see and on which date and how many tickets

you want. The agent will tell you what seats are available. You might ask what the prices are. The agent will give you prices for the various seating areas. You'll say which seats you want. The agent will ask how you're going to pay. You'll say you'll use your credit card and give the number and the expiry date. You'll probably give your address, too, so the tickets can be mailed out. Then you'll both say thank you and hang up, and you'll stand there looking at the price you jotted down and wondering if two hours of Leonard Slatkin conducting the Washington National Symphony Orchestra can possibly be worth that much.

There are clearly some rules governing how this interaction proceeds. You can't ask the price before you specify the show, for example. If the only available seats are in the stalls, you can't ask for circle tickets, and so on. If you thought about it a bit, you could probably sit down and write a set of rules somehow, or maybe even invent a diagram, to describe the valid sequence of events that should be followed during the purchase of theatre tickets. What you'd be writing is a protocol. Computer protocols are very similar. They describe the valid commands, queries and responses throughout any given process or transaction between two computers. Computer protocols also usually define how the computers are to represent the commands and data they are sharing. In our analogy, this might be equivalent to specifying that the conversation with the ticket agent must be held in English.

A protocol, then, means a pre-defined and formal set of rules which define how communication takes place between two or more computers.

Protocol Stacks

Protocols are often used in combination, and most often they are 'layered' one above the other. Typically, the lower layers provide some sort of supporting role in helping the higher layers to do their job. In our theatre booking, there is an underlying protocol operated by the telephone system to establish a connection between the phones so that each party can hear the other. So we have the 'ticket booking' protocol layered over the 'telephone connecting protocol'. In computer jargon, these protocol layers are usually called 'stacks', and they're what makes the Internet work.

OK, back to the plot. Our web browser now knows it needs to contact the web server on the machine called **www.linuxformat.co.uk**. It's going to use the stack of protocols shown in **Fig2**: namely, http (HyperText Transfer Protocol) over

SO MANY Ws, SO LITTLE TIME...

We don't know about you, but this convention of putting 'www' at the start of the names of web servers is starting to really annoy. Yes, it stands for 'world wide web', but the fact remains that in the English alphabet, 'w' is the only letter which requires three syllables to pronounce – in fact it's the only non-monosyllabic letter in the whole alphabet – unless you live in Texas, in which case 'Dubya' is still a two-syllable pronunciation.

It's bad enough putting one 'w' at the start of the name, but three? Why can't we have addresses like web.somecompany.co.uk, which would do just as well and wouldn't need you to make a noise like a galloping horse to pronounce it?

the Transmission Control Protocol (TCP) over the Internet Protocol (IP). Let's look at what each of these protocols does. It's easier to explain if we start at the bottom of the stack, with IP.

Down on the Wire

IP is responsible for routing packets of data (they are called 'IP datagrams') around the Internet to reach the correct computer.

Imagine a small local area network of computers connected via a single network as shown in **Fig3**. We'll assume that the network is an Ethernet but the same principles apply to other LAN technologies. Each computer has an interface card (sometimes called a network adaptor or NIC) which connects it to the network. Built-in to this card is a fixed 'address' (actually it's a number) which identifies the card. This is called the Ethernet address or the MAC address. Any computer can send data to any other machine on the same network by sending out a packet (more technically called an Ethernet frame) which has the MAC address of the intended recipient at the beginning. The network interfaces of every computer on this network will see this frame as it passes by, but only one (the one whose MAC address matches the one in the packet) will actually capture it.

In reality, the world's computers aren't connected together onto a single cable. They are built as a set of inter-connected networks, or Internets. **Fig5** shows how this is done. Now things start to get more complicated and interesting. Suppose machine **B** wants to communicate with machine **Q**. It's no good **B** hopefully launching an Ethernet frame onto its network, because **Q** won't see it and the frame will simply be ignored. There needs to be some way for the system to figure out that to get from **B** to **Q** you have to go through **G**. This is called routing. Ethernet can't do this – it's not a routable protocol. This is where IP comes in.

IP – The unreliable postman

Every computer connected to the Internet is assigned an IP address, which is a 32-bit number. These are written in a rather quirky way called dotted decimal notation. You take each 8-bit piece of the address, write it as an ordinary decimal number, and put dots between them. For example, the binary IP address 10111000000001100001001100010111 is written as 184.6.19.23. You have a choice – work it out for yourself, or trust us!

Now comes the really clever part. The IP address is divided into two parts. The top part, called the network ID, identifies the network (inevitably) that the machine is on, and the bottom part, called the host ID, similarly identifies a specific host on that network. For reasons which remain a mystery to many of us, a computer connected to a TCP/IP network is called a host.

The network ID is used to route the data to the correct network, and the host ID is used to select the correct machine on that network to receive the data. The boundary between the network ID and the host ID can in fact fall anywhere within the 32 bits, but to keep things simple we're only going to consider the (common) case where there are 24 bits of network ID and 8 bits of host ID. This boundary corresponds to having the first three numbers of the dotted decimal address being the network ID, and the bottom number being the host ID. So in our example, the network ID is 184.6.19 and the host ID (on that network) is **23**. See **Fig4**.

OK, so how does IP work? Let's take the simplest case of a machine communicating with another machine on the same network – say machine **B** to machine **C** in **Fig5**. The IP layer recognises that the destination address (184.6.19.24) is on its own network, so it can




Fig 1. The components of a URL




Fig2 Protocol stacks and encapsulation.




Fig3 A simple network.

send the data directly to machine **C**'s MAC address, if only it knew what it was. There is a protocol called the Address Resolution Protocol (ARP) which is used to discover a machine's MAC address if its IP address is known. ARP works by broadcasting a request onto the local network, asking "what's the MAC address for the machine with IP address 184.6.19.24?". Since the request is a broadcast it is received by all the machines on the network, but only machine **C** recognises the target IP address as its own, and replies direct to **B**, saying "that's me, and my MAC address is such-and-such". Now **B** can build an Ethernet frame with the correct destination address to send the data to machine **C**. If you're thinking it seems a bit inefficient if you have to broadcast an ARP request every time you want to send data, you'd be right. In fact, what happens is that machines remember the MAC addresses of other machines they have recently contacted. This store of known MAC addresses is known as the ARP cache.

Routing

Now let's return to our example of machine **B** sending data to machine **Q**. Since **Q**'s network ID (184.6.20) doesn't match **B**'s network ID, the IP layer on machine **B** consults an internal



TUTORIAL TCP/IP




Fig4 The components of an IP address.




Fig5 An example of small intranet.

◀ routing table to figure out what to do. Routing tables contain entries of the form:

To get to network X, send the data via IP address Y

In the case of machine **B**, the routing table is very simple, consisting of a single default route which says “if you don’t have any more explicit instructions on how to route this data, send it to address Z”. In this case, the default route goes to one of machine **G**’s network interfaces, 184.6.19.200 (in fact, the great majority of machines on the Internet have only “one way out” of the local network, and therefore need only a default route). Now, machine **G** has connections to two networks. In traditional UNIX TCP/IP vocabulary, **G** is called a Gateway. Elsewhere in the network community it’s called a router. Notice that each of **G**’s network interfaces has its own IP address – one on each network.

On with the story. **B** will construct an IP datagram with machine **Q**’s IP address in the header, but it will actually send it to **G** (of course, first time around **B** will have to generate an ARP request to find **G**’s MAC address).

G will examine the destination address of the datagram and realise that the machine it’s intended for is on a network that it’s directly connected to (184.6.20) and so will use ARP to find **Q**’s

MAC address, re-encapsulate the IP datagram inside a new ethernet frame, and deliver it to **Q**. Notice that the datagram gets a new ethernet frame with a new MAC address at each station along its journey, but the destination IP address stays the same throughout the process.

The process can get more complicated – for example, if **G** receives a datagram from **B** which is destined for some other network, **G** will have to forward it via its own default gateway, **X**. But hopefully you have the general idea. The point to remember is this: the IP protocol is capable of routing a datagram across the Internet based on its destination IP address.

Naming Services

Next time you’re online, try typing the IP address 212.113.202.71 into your browser. It’s the IP address of *Linux Format*’s web server (or at least it is as I’m writing this). It should work. However, it’s not the most friendly way of identifying a machine. Most people don’t want to remember a stream of numbers just to log check out a website, they want to work with names which are easier to remember. To this end the Internet has evolved a naming scheme, with names like www.linuxformat.co.uk standing in for numerical IP addresses. When you type such a name into a browser, the browser has to figure out what the corresponding IP address is. This process is called ‘name resolution’.

In the early days of the Internet, a centralised ‘hosts’ file was kept on one specific machine with the name-to-address mappings of all known machines. The file was copied out to every machine on a regular basis. But as the Net grew, the file got too big, and the frequency of updates to it became unworkable. And so a distributed name service called DNS was created. ‘Distributed’ means that knowledge of the names and addresses of all the machines on the Internet is spread across many DNS name servers. It is many years since any one computer knew about the whole Internet.

Each name server knows about the names and addresses within one particular section – or domain – of the hierarchical namespace which DNS defines. If you want to know what a ‘hierarchical namespace’ is, you’ll have to read the rest of this series. In the meantime, it’s a wonderful phrase with which to bluff your way in computing. (“Oh, I spent the weekend cleaning out our hierarchical namespace – it was absolutely filthy!”) It is difficult to over-emphasise the importance of DNS in holding the Internet together. It is a vital piece of network infrastructure.

The Transmission Control Protocol

To continue with our story: IP cannot perform the whole job of delivering data to the web server. There are a couple of key things it doesn’t do. First, it’s a ‘best effort’ rather than a guaranteed delivery service. Sometimes, if the gateways run out of memory

NAMES vs ADDRESSES

Stack levels

As we talk about translating between names and addresses, let’s be clear about the distinction we’re making when it comes to computers. In normal English usage, your name identifies you as an individual and your address identifies the house, street and city where you live. A machine’s name (such as www.linuxformat.co.uk) and its IP address (and for that matter its MAC address) are simply different ways to identify a machine that are used at different levels in the stack.

STANDARD PORT NUMBERS

What the servers are listening out for...

daytime	13	bootpc	68
ftp	21	http	80
ssh	22	pop3	110
telnet	23	sunrpc	111
smtp	25	nntp	119
dns	53	irc	194
bootps	67	mysql	3306

to store incoming datagrams for instance, IP traffic is lost. Secondly, IP only does machine to machine delivery. It has no notion of delivering data to a specific service within the machine. For example, it can't direct data specifically to a web server, or a telnet (remote login) server, or a print server. These facilities are provided by a protocol called TCP, or Transmission Control Protocol which is layered above IP.

TCP delivers data to a specific transport endpoint, also known as a socket, which is identified by a port number. Servers are said to "listen for connection requests" on specific ports, and there is a fixed association between the services and the ports they use. For example, web servers always listen on port 80, ftp servers on port 21, and the DNS servers we talked about earlier on port 53. If a DNS server woke up one morning and decided to listen on port 54 for a change, it would have a very quiet day because no-one would know where to find it.

In addition to delivering data to a specific transport endpoint based on the port number, TCP also creates the illusion of a reliable connection between the client and the server. This is a pretty neat trick, given that it's using the non-guaranteed IP delivery service to deliver the data. The way it does this is both interesting and elegant but we don't have enough space to describe it in the required detail here. Basically, each chunk of data which TCP sends (it's actually called a segment in TCP parlance) is labelled with a sequence number when it's sent out. The recipient acknowledges receipt by sending back 'ACK' packets containing an acknowledgement number. This allows the TCP at the sending end to figure out if a segment has gone missing. If so, it re-sends it.

Fig6 shows how the combination of the server's "well-known" port number, combined with the IP address of the machine where the server is running, identifies precisely where the client needs to connect to. The client's machine has an IP address as well, of course, and in fact the client will allocate a port number to its own end of the connection. (That's just the way TCP works). However, unlike the port number at the server end, the client's port is essentially arbitrary – it will likely pick a different one each time. It is the association of four items:

- Client port number
- Client IP address
- Server port number
- Server IP address

These uniquely define the connection. This is an important point to remember if you come to start configuring firewalls because they usually base decisions on whether to forward packets on these four values.

Fig6 also shows a couple of other servers listening for connections on their own well-known port numbers.




Fig6 Association between client and server.

The Application Protocol

Application protocols are generally to be found at the top of the protocol stack. They perform some application-specific task such as serving web pages or allowing data to be spooled to a remote printer. All the underlying protocols only exist to allow the application protocols to do their work. In the case of our web browser the application protocol is http, but it could just as easily refer to ftp, https and smtp. In the case of http, the protocol allows browsers to request pages from web servers, passes in any form data which the user may have entered, and passes HTML pages back.

Encapsulation

Fig 2 summarises the whole process, showing how an http request proceeds down the protocol stack at the sending end. As it passes down to each layer, that layer adds its own 'header' information to allow it to do its work. Each layer's header becomes part of the 'payload' of the layer below. This process is called encapsulation.

After the resulting IP datagram has made its way to the destination machine, it passes back up through the protocol layers. As it does so, each header is stripped away once it has performed its function. Each layer is not aware of the header added by the layer below to enable it to do its job. Up at the application layer, the web server simply receives the http request packet exactly as the browser created it. Conversely, the lower layers have no idea of the meaning of the data they are carrying as a payload.

There's a good real-world analogy to encapsulation, if you can remember as far back as the days when people used to write real letters with real pens on real pieces of paper. I might write a letter to a business colleague, John Smith say, and hand it to a secretary (a lower layer of the protocol stack) saying "send this to John Smith". The secretary will put the letter into an envelope, write John's address on the front, and post it. Next day the letter is received by John's secretary, who takes it out of the envelope, throws the envelope away, and hands the letter to John. The two important things to note are that neither John nor I ever sees the envelope, and neither secretary needs to read or be able to understand the content of the letter. That's encapsulation in a nutshell, so to speak. [LXF](#)

NEXT MONTH

We've seen briefly how the Internet works and we've identified a couple of key components – routers and name servers – which are needed for it to function. In the next part of the series we'll look at how Linux can provide those components and also how it can provide a (regrettably) essential component of every company's network – the firewall.

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



Damn(small) Linux!

Q I have to make a plea; despite attempting many times, I still am not able to install DamnSmallLinux (from LXF41 or LXF42, I've tried both).


I have access to an Athlon 1100 with CD Writer (Nero 5) & DVD, and I get the DVD subscription edition so that I don't miss out on the interesting distros that you feature occasionally. Why I'm particularly interested in DSL is that I think it would be very appropriate for my computer (hardware: An 7/8 year old P120 with a 4GB HD, a second HD of 203MB, running [now] Win 95c). The second HD is where the DSL would be, obviously.

I have tried to follow all the instructions, including those in LXF41, p105 (Creating the ISO with Windows; at least on the Athlon I've got access to, running Win2KPro *winnkiso* is not recognised).

I really, really would like to install DamnSmallLinux on my computer – and more generally to learn about how to install distros – I'm afraid possibly in it would need to be in words of one syllable, however. I am possibly revealing my irredeemable stupidity but I need a hand up!

Vic White, via email

A You should be able to make yourself an ISO using *Nero*, rather than *winnkiso*, simply by pointing it to the appropriate location on the DVD, then burning it to CD. Once you've got the CD, you can boot off it and install DSL as normal. Everything will be nice and easy drag-and-drop, at least on the *Nero* side of things, and you can get DSL up and running without a whole lot of difficult messing. Then you'll be thrust head first into Linux with a minimal set of GUI applications, so save up your command line interface time for that!



DamnSmallLinux is a cut-down Knoppix: great for low-end systems.

Debugging cron job

Q I've created a simple script (*chmod +x*) that creates a couple of text files, and ftp's one of them to an ftp server. The script runs fine via the command-line.

If I use my *crontab* to run the script (every day) all the text files get created OK (so *crontab* works) but resulting text file does not get sent via ftp. I can't find any error messages. Here's the *crontab* entry:

```
00 09 *** /home/david/cronstuff/
uploadIP
```

And here's the failing (final) line in the text file 'uploadIP'

```
/usr/bin/ftp -i < /D/home/david/
Desktop/crontab/ftp.txt
ftp.ftp site.com
```

...and ftp.txt is my login script containing **put uploadIP**. As I said earlier, it all works fine when run via the command-line. The problem seems to be *crontab* will not run

the ftp command.

Crontab is run as the user it belongs to (me) and I've got my *.netrc* script configured correctly – could it be that *crontab* is NOT using the *.netrc* file? Please help!

From the LXF forums.

A The first step in debugging any *cron* job is to run from *cron* and see exactly what it does by piping the output somewhere useful. You are referring to both */home/david* as well as */D/home/david*, so paths need to be accurate. When a



A *cron* job is generally simple to create, but one should remember to define the *PATH* variable.

cron job produces an error, it will email it to the user who is running the job, so checking the contents of \$MAIL for useful errors can always be helpful. You need to define some basic variables in your *crontab* file, such as HOME to point to /home/david, PATH, which should contain your search path for binaries, and anything else which is needed by the processes you are using.

If you can run the script from *cron* and provide us with the error output, if any, or post them to the LXF forums, then we can help you further in debugging your script.


Remote access

Q At last BT have broadband enabled my local exchange. I have purchased a D-Link DSL-604+ wireless ADSL Router. This works fine (although it does get a bit hot, (but with winter upon us, it doubles as an underpants warmer for these cold mornings). At last I can surf in Linux from

my wired PC, no more WinModem headaches, and I can also surf from the garden with my WiFi laptop. I would now like to be able to access my home PC to transfer files, have a peek at my webcam etc from the office PC – but how?

The router is set for IP 192.168.0.1 and it is set to assign IP addresses to any PC on my local home network. However BT assign my router with a Dynamic IP address. (I was told that the ADSL service does not provide static IP addresses.) On the web I have seen services advertised, such as no-ip.com, to provide redirection for people with dynamic addressing. Is this the only way to get remote access to my home network?

Also I have an old mini-tower that I could press into use as a web server and have a little play around with *Apache* and the like. How would I get its name to be published on the Internet? As a subscriber to LXF since issue one, thanks for all



Many DSL providers block port 80, but help is at hand!

your past help. You really do an excellent job with the mag. I hope you can help me with my problem. *Melvyn Pearce, via email*

A Services such as dyndns.org and no-ip.com allow you to change a hostname assigned

to a system on a dynamic IP address. So, if you have 'bob.no-ip.com' assigned to your system, as BT switch your IP address, the hostname will update accordingly. However, this is only one piece of the puzzle. If you want to run a web server off that system, you need to first find out if BT permit you to run a web server as part of your service, as many DSL providers will block port 80, plus you will need to setup port forwarding on the router. If you choose to obtain a domain name, rather than use the no-ip.com or dyndns.org hostname, it is trivial to CNAME the domain onto the dynamic hostname.

Remote access can be achieved by port forwarding *SSH* through the router, either running it on port 22 if that is not blocked by BT, or running it on an obscure high port. Once you have *SSH* available remotely, then you can port forward other services, such as *X* or *VNC*.

Red Hat? Nein!

Q I have finally got around to setting up Red Hat Linux 9 on my PC but I have a couple of problems and I don't know where I should go to find the answers I need:

1 I have a mini-network running at home under Windows. How do I set up my network card to connect to it? When I look at the info on the card I have the following error:

Cannot activate network device eth0
The reason given is:
Determining IP information for eth0 failed.

The settings are as follows:


Hardware: device VIA VT86c100A ➤➤

A QUICK REFERENCE TO: Snort

A firewall isn't the last line of defence against intruders, so here's the low-down on IDS...

Even with a comprehensive firewall configuration, there is still the problem that we need to have at least a few basic ports open so people can access things on our network. Whether this is limited to *SSH* and *Apache*, or includes more complex and vulnerable processes, including *Sendmail* or *SNMP*. We can't simply block all access, however there are often exploits for CGI or PHP scripts, or someone can attempt certain exploits on a service, such as buffer overflows on *Sendmail* or *CodeRed* against *Apache* thinking it's *IIS*.

Monitoring traffic permitted through the firewall is the job of an Intrusion Detection System, (IDS), which continually monitors network traffic and performs packet analysis to look for signatures from viral activity, exploits, or other unwanted happenings on our network. *Snort* is a wonderful Open Source IDS, and coupled with a Web interface such as *ACID*, or using a notification



Snort is a wonderful addition to any secure network, so long as the ruleset is tuned to avoid false positives.

systems similar to *PigSentry* to tell us when something unwanted happens.

Snort can be quite a pain to set up, though many prebuilt signatures and rules are available online, so we don't always have to start from scratch. It's important to ensure that the rules are applicable to the environment, so if there isn't an *IIS* server, then we can just skip all of the *CodeRed* nonsense. Anyone who gets notified when there is an IDS event on a web server will really be unhappy if they get told someone

tired *CodeRed* or *Nimda* against *Apache* or *Boa*. However, other attacks, such as PHP exploits or file insertion exploits through CGI are invaluable when securing a network from the inside out.

Snort can be downloaded from www.snort.org, although most distributions will ship with *Snort*. There are some great docs on building a ready-to-roll IDS, such as the *Snort* and *ACID* HOW-TO at www.floridahoneynet.org/whites/snortacid.html

FREQUENTLY ASKED QUESTIONS LINUX AND THE INTERNET

FAQ HOW DO I CONNECT TO THE INTERNET WITH A MODEM UNDER LINUX?

There are many different ways to set up a dialup connection with Linux. Assuming we're after a standard PPP connection, then we can use something like *kppp*, *gnome-ppp*, *wvdial*, or one of the numerous distribution-specific diallers. With something as basic as *kppp*, you simply enter the appropriate phone number, along with your user name and password, then try it out and see what happens. Of course, we can also do it by hand, by creating chap scripts and creating a `/etc/ppp/` options for when we ppp-on:

```
/etc/ppp/options;
/dev/ttyS0
crtscts
modem
lock
connect 'chat -f /etc/ppp/chat-script'
defaultroute
mru 576
mtu 576

/etc/ppp/chat-script;
ABORT 'NO CARRIER'
ABORT 'BUSY'
" ATZ
```



Linux supports modems better these days.

OK ATDT08452121666

ogin: username

word: password

ocol: ppp

FAQ WHERE IS COM1 IN LINUX?

Linux uses `/dev/ttySx` for serial ports, and COM1 is referred to as `/dev/ttyS0`. You can check to make sure your modem is really connected to `/dev/ttyS0` with a program such as *minicom*, where you can send ATZ to the modem and check to ensure that it responds with **OK**.

FAQ WHAT ARE PAP AND CHAP?

PAP and CHAP are different authentication methods used by PPP

servers, and which one you need to use will depend upon your ISP. Nearly all ISPs use PAP, although users of BT Openworld's dialup services will need to use CHAP.

FAQ WILL MY WINMODEM WORK?

Depending upon the chipset used by the WinModem, it may or may not work. Lucent models generally work better than others, in so far as that they actually occasionally work at all. www.linmodems.org has docs, and a lot of links to kernel modules for a number of common chip sets, although your mileage may vary.

FAQ HOW DO I USE MY CABLE MODEM?

Cable modems are generally connected to a NIC on the computer, so we need to use DHCP, rather than PPP, to connect to the Internet. There are many DHCP clients available, including *dhcpcd*, *dhclient* and *pump*, although generally a distribution will only have one of them installed. *pump* is usually the most successful with cable modems, so we can configure our NIC with:

```
# pump -i eth0
```

For a more permanent fix, you will need to reconfigure your networking within your distribution, in order for the DHCP client to be started at boot time, rather than having to do it by hand all of the time. If you have limited success with *pump*, then it's often best to try another client, such as *dhclient*, as some don't like working with certain DHCP servers.

◀◀ **Rhine II PCI Ethernet (this is an onboard net connection not a separate PCI card). DNS: Hostname set to localhost.localdomain (all other settings are empty) Hosts: IP 127.0.0.1 Name Localhost.localdomain aliases localhost**
I presume that *localhost* and *localhost* have to be renamed to match my network but I am unsure as to what goes where.

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!

2 My second problem is setting up the Internet under Red Hat. I have a BT Voyager ADSL modem (which actually comes up as Globespan USB ADSL WAN modem under Windows). I connect to the Net using a BT broadband connection. How do I set up the modem under Linux and what info do I need from BT to make the connection?

I know there are places on the Net where I can obtain this info but I haven't got a clue where to look. If you could point me in the right direction or help me with the answers, I would be much obliged!
Antony Wright, via email

A As you did not supply us with any information as to the IP configuration of your network, we can only make assumptions as to the problem. It would appear that eth0 is attempting to use DHCP to

acquire the network configuration required, however it is failing to do so. If you do not have a functional DHCP server available on your network, then it will simply not be able to do anything. You should be able to duplicate the IP configuration from Windows and manually enter it into Linux to gain access to your network. You may want to update `/etc/hosts` with the IPs and hostnames of the systems on the network, although you can still access them via IP without updating it.

There are up-to-date drivers available for your USB DSL modem, which is available from <http://eciadsl.flashhtux.org/>. The site itself is in French, but using BabelFish translation and/or Google should help you locate the relevant documentation that you need in English from various mailing lists.

More rsync about

Q I just picked up a copy of your magazine in the US, and wondered if you might be able to help me out with a little issue. I have a bunch of websites hosted on a UserModeLinux host (www.linode.com) and I'm trying to make sure I have them all backed up properly. My backup script is basically as follows:

```
tar -zcvpf /var/backup/virtual.site.tgz
--directory=/var/www/html/dir .
```

It works fine when I run it as root from the command line, but when run as a *cron* job (as root) it chokes on some JPEG files. When I *untar/unzip* the file, it says **'unexpected end of file'**. Any ideas? Are the JPEGs throwing it off somehow, and if so, what's my way around this?

Tom Haddon, via email

FAQ WILL MY FROG WORK WITH LINUX?

Unfortunately under Linux, frogs are limited to 2400baud, and that is assuming you can make them sit still on their lily pads long enough to negotiate the PPP connection. Often, an upgrade to a toad or newt will solve the problem.

If you're actually referring to the Alcatel SpeedTouch USB DSL modem, used by the majority of ISPs in the UK, then it will work using the binary kernel modules distributed by Alcatel. You will need to compile *PPP-over-ATM* (*pppoa*) into the kernel, as well as obtain a number of user-space configuration programs, but many people do use this system for their Internet access. A fairly well-written and comprehensive HOWTO is at:

<http://linux-usb.sourceforge.net/SpeedTouch/howto.html>

FAQ CAN I SHARE MY INTERNET CONNECTION BETWEEN BOTH OF MY SYSTEMS?

Assuming you already have a functional network using private IP addresses, it's very simple to allow other machines to use your Internet connection. If you're running a 2.2.x kernel on the machine with the modem, then you will need to recompile the kernel with *ipchains* support, and install the *ipchains* package. For 2.4.x users, there is the choice of *iptables* or *ipchains*, although

the former should be chosen if you don't already have an *ipchains* configuration, or have the time to start again from scratch. We can share our connection by using IP Masquerading, which is done with one of the following commands depending on your choice of firewall:

```
# ipchains -A FORWARD -j MASQ
```

or

```
# iptables -t nat -A POSTROUTING -o ppp0 -j MASQUERADE
```

```
# modprobe ip_conntrack_ftp
```

We also need to enable IP forwarding, which is done with:

```
# echo 1 > /proc/sys/net/ipv4/ip_forward
```

FAQ HOW DO I SECURE MY SYSTEM AGAINST ATTACKERS?

Firewalling external connections is easily done with *iptables*, assuming that you don't want to permit any incoming connections:

```
# iptables -A INPUT -i eth0 -j ACCEPT
# iptables -A INPUT -i ppp0 -m state --state ESTABLISHED, RELATED -j ACCEPT
# iptables -A INPUT -i ppp0 -j LOG
# iptables -A INPUT -i ppp0 -j DROP
```

A When a *cron* job exits, it will email the output from the command to the user who ran it. If your tar command dies at some point during execution, then the error will be emailed to the root user. Without this error, it's somewhat difficult to put a specific reason to why it is not working.

If you are intending to backup the site to an off-server location, then the *rsync* tool is more appropriate for this, as it only requires the transfer of files which have changed, rather than shovel the whole tarball over the Internet. See *LXF44's What On Earth* (page 52) for more information about using *rsync*.

Wasted effort?

Q Would there be any point in me trying to compile source code on a P200MMX, 98MB

RAM, 2GB hard drive, Matrox 2064 Millennium? The times would be long, but would it be worthwhile, what with PCs today being as fast as they are?

Keith Scott, via email

A Unless there is a special reason to build code from source, it's probably not worth the effort on such a slow system. As nearly everything is available in a binary package of one form or another, you're going to be waiting a rather long time to get anything done on a system as underpowered as this.

There's nothing stopping you, other than the low amount of RAM, but really, why would you want to?

Three problems

Q I loaded Red Hat 9 (dual booting with WinXP) on my home desktop, but the GUI



nVIDIA's binary Linux X drivers – its video cards can work with XFree86.

did not come up, have also got the drivers from nVIDIA, but do not know how to load it in.

2 What editor can I use and the command to type to get it out if I need to edit some config files from the Linux prompt. I can't get the GUI up (due to problem number 1)

3 How come the eth0 does not come up (and give me an IP, like Windows boxes) when I boot up the Linux box? I'm trying to plug it into a switch, connecting it to my other boxes (they are Windows-based, unfortunately). The Windows machines autoconfig and give themselves an IP, so I was wondering what to do to get my Linux box to do the same (without having to manually give it an IP every time). Thanks!

Roy-Yew-Wah Leong, via email

A The nVIDIA site provides documentation on setting up the XF86Config file for use with its video card drivers, although it does require some basic understanding of how Unix works, such as path names. The documentation that comes with the tarball should be fairly simple to follow, or you can at least get X up and running using the *nv* driver and work from there.

Both *Vi* and *Emacs* are console-based text editors which can be used without X, although they take a little getting used to. *Emacs* is generally a little simpler to use from the outset, although there are numerous other alternatives, such as *vim*, *jed* and *pico*. As you are a newbie with Linux, a beginners book on general Unix or

Linux is probably a valuable investment. Linux will either allow the NIC to automatically configure using DHCP if you have a network router, or you can supply manual configuration information, which persists across reboots. There should be no need to manually enter IP information every time you reboot into Linux.

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



Firewall-starter

Q I would like to follow up *Hardware firewall* in LXF45. I am running two machines, one a Windows machine, and the other a SUSE 8.2 – *Apache* web server. I have DSL and I am using a *LinkSys 8* port router for port redirection. The ftp, telnet and http ports are redirected to the SUSE machine. No other ports are redirected. I have ‘upgraded’ the router’s firmware with the most updated version. Would it be too hard for ‘script kiddies’ to attack and/or crack this type of setup? Thanks for your advice and input and I love your magazine!

Burt, via email

A Burt, if your hardware is always flashed to the newest version of the firmware, then you are fairly safe from the casual ‘script kiddie’ getting into your system via a known vulnerability. I’m a little concerned that you have opened up telnet on your SUSE system. Although there are no issues with telnet from a code point of view, the protocol is flawed in that it sends password and data across the network in clear text. Any security issues you may experience are likely to be related to the services available publicly on your Suse system rather than the firewall.

iptables rules, OK?

Q I’m a Red Hat 8.0 user and I have just downloaded the newest version of *iptables* via the Red Hat Network. It was a masked package but I was able to force it to install anyway. That was my first mistake!

Once installed, I copied my rules back into place from my previous version’s configuration. Now, when I restart *iptables* it complains about all the entries, which use connection states. Below is an example of a command I am trying to pass:

```
iptables -A INPUT -m state --state INVALID -j DROP
```

When I run this command, *iptables* tells me that there is no such table or rule. I’ve even tried to downgrade my Red Hat system back to the old version (1.2.6), but it’s still doing it. How can I secure my server while still allowing stateful filtering on my system?
Terry, via email

A I ran into this exact problem on a customer’s server and I’ll admit that it made me pull my hair out for a while. You need to be sure that you have cleared out the old *ip_contrack* modules before trying to start the newer version of *iptables*. All the state arguments require this module to function properly. You should be able to *rmmmod* the old modules once you’ve stopped the *iptables* service and they should load up automatically once the new *iptables* fires up. If you’re paranoid, you

might want to bring down your public Ethernet interface before bringing down *iptables* or just do a total system restart after the RPM installation.

Heading

Q The company I work for is pretty much entirely a Microsoft shop. I’ve been trying very hard to get some Linux boxes up and running, and it seems that people are starting to listen. My boss has let me set up a DHCP server and he’s disabled DHCP on the Windows 2003 server. I guess he thinks it’s easy enough to get it back if the Linux box dies... It’s all been going well for about a month and I’ve hardly even had to log into the server, but the other day I did though, just to see how it was doing.

The first thing I noticed was that the clock was about 3 hours off. This is nothing serious, considering how old the hardware is, but I know my boss will try to blame it on the ‘flaky’ operating system I installed. Is there anything you can recommend to make the system ignore the hardware clock and trust its own clock? Thanks for the advice!
Adrian, via email

A First off, Congratulations on getting Tux’s foot in the door at your company!

Fortunately, there are plenty of time-related tools in Linux to resolve your issue. It is not possible to set the system clock to run independently of the hardware clock, but there are ways to modify it.

To begin, I would recommend setting up *NTP* on your server. *NTP* is a protocol for getting the current time from designated time servers on the Internet. Many of these servers are attached to atomic clocks for the utmost accuracy. Once it is installed, all you’ll need to do is find an *NTP* server that is not too far away from you and set it to go.


There is also a useful utility called *hwclock*, which allows you to manipulate the hardware clock. Combined with *NTP* this should be all you need to alleviate the systems of a dodgy hardware clock. The command: `hwclock --systohc` sets the hardware clock to the current time in the system clock without having to reboot and enter the BIOS. Coupled with *NTP* and run every few hours in a *cron* job, you shouldn’t see the hardware clock drift.

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 the company website at www.rackspace.co.uk



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 **Andrew** – your new Zaurus PDA will be with you shortly!

STAR Q

Q Over the course of many years, our geek trio has collected several computers and now that our main systems are setup, we don't really use the older Pentium 2s and 486 machines we have lying about for much except playing with. This is where our question begins...

What we would like to do is create a cluster out of them. We want the cluster to run dedicated servers for web serving, mail serving, game serving, distributed computing and number-crunching. The ideal usage scenario would be getting compilation of our programmes to be done on the cluster, alleviating our main machines of the load for other tasks. We've investigated a few different means of doing this and

the two main methods to achieve our goals seem to be either OpenMosix or Beowulf. What method would be best, and how would we go about implementing it?

Andrew, via email

A Thank you for your interesting question. As you have already mentioned, there are two well-known and well-adopted clustering options for Linux. They work very differently however, and you'll probably find that only one is suitable to your needs. In a nutshell, Beowulf is more established and extremely powerful, but it does require the application that you are planning on using to use to be cluster-aware. Many high-end computing applications such as renderers (and even some compilers) are Beowulf-aware. To ascertain if an application will work on your Beowulf cluster, you should check that it uses

MPI or PVM clustering libraries. Chances are, however, that the generic programmes you want to run aren't ready for this. A more likely candidate for your proposed geek-cluster would be OpenMosix.

OpenMosix makes use of a customised Linux kernel to handle the job of passing processes between nodes in the cluster. This kernel keeps track of the best available node and will pass processes to this node. It is because of the fact that this functionality is built into the kernel that any process can be passed to another node (very similar to processes being passed between CPUs on a multiprocessor system). This is ideal for your situation – as a test, you can start up four instances of SETI@home on your client, and the tasks would be offloaded to the four best nodes. Similarly, when compiling, each process of the

compilation will be able to run on a separate physical processor. Remember that not all programmes play nicely when trying to compile them on multiple CPUs.

As far as implementation goes, you should look at nothing less than a 100Mb/s network, preferably switched. In brief, you will need to apply the OpenMosix kernel patch to a stock 2.4 kernel source and compile (although Red Hat and Debian packages are available too). You'll also need to compile the *userland* utilities. With these two steps done, you can get down to the actual configuration. I'm intentionally avoiding going into the actual configuration because I'd only be plagiarising the excellent documentation on the openMosix site <http://openmosix.sourceforge.net/>. Take a look there and if you run into any hiccups – or just want to brag – drop us another email.

Sendmail response

Q I am running *Sendmail* and it appears as though I am getting an ever-increasing queue size. When I run `sendmail -bp` it lists thousands of domains, that all exist. When I do a `sendmail -q` it seems to take an eternity to run through the queue and it

always says, '**Cannot establish an SMTP connection**'. I can however telnet to the remote server on port 25. What I would like to know is how I can turn the *Sendmail* debug up so that I can see *Sendmail's* server response.

A Generally, to debug a SMTP connection, you would use the `sendmail -X` flag. Here is an

extract from the man file on debug.

`-X logfile`

Log all traffic in and out of mailers in the indicated log file. This should only be used as a last resort for debugging mailer bugs. It will log a lot of data very quickly.

On a busy mail server (which it sounds like you have!) this will increase the workload by a fair amount. I would try

to be as careful as possible using this command. Also bear in mind that some versions of *Sendmail* require you to specify a to: address after the filename, for example:

```
sendmail -X filename
user@domain.com
```

I hope that this helps resolve your issues or at least helps you find and fix the root of the problem.

« Printing & SUSE

Q I use *KMail* routinely as my mail client. I run a SUSE 8.2 system on an Athlon 1.2GB with 512MB RAM, and KDE is my standard window manager. There are no serious modifications, although I run Win4Lin v5 and usually work with the Win4Lin kernel as I never know when my work will require moving temporarily into Win4Lin. I have downloaded and installed the TrueType fonts from SUSE's site and installed some extra ones I need via the *KDE Control Center*.

Occasionally, I print an email directly from within *KMail*. Up to now, it has been a matter of selecting the email I want to print, right-clicking and selecting print and away it goes; or select the print icon on the tool bar and similarly it gets printed. Now it won't.

Curiously, everything else still works. I can print from within Win4Lin, *Open Office1.1* still works, I can save the *KMail* file and then open and amend and print the bits I want from within *KWrite*, but I cannot print from within *KMail*. It is almost as if a link has been broken but I have no idea where. I can't be very helpful on this one, as there is really nothing complex that I know how to analyse, but the attempt below might put you on the track. Whatever is happening is at system level I think, and although I do not know this system awfully well, my impression is that it is somewhere around the print spooler and how it sends the info to the printer, but I have no idea where to look.

I have tried to do some tracking down via the *KJobViewer* and found that a previous one sent by *KMail* which did work, registered one page printed. Those requests which do not work show 0 page printed – as I guess you would expect. However, when I right-click on one of the previous jobs sent and open its attributes, the bottom line of 'job-state-reasons' says: 'aborted-by-system'. A job which did get printed successfully says: 'job-completed-successfully'.

Logical, isn't it? However, where on earth do I look next? I suspect that this is some silly small thing that has been broken in some way, but I have no idea where to go. Any suggestions, remembering I am not

a good command-line operator – but I am willing to try as long as it is not too technical and requires a reasonably simple procedure? Also, please, please, please *don't* tell me to use *Vi* or *Vim* – I can't. But I can use *KWrite*. Thanks!

Dr Tony Young, via email

A Linux printing generally allows one to supply a variety of options to the *lpr* command, such as the particular printer to use, the number of copies, and any other specific options pertaining to the printing. Should this command-line become confused, then the process sending the task to the printer will bail out. As everything else works, yet *KMail* does not, it would suggest a problem specific to *KMail*, rather than a concern with the *lpd* or *CUPS* configuration on the system.

If you verify the configuration information when printing from *KMail* and compare it to another KDE application, it may be that *KMail* is trying to print to a printer which no longer exists, or has a messed up command line somewhere.

Man... Drake?

Q I have just installed MDK 9.2 (thanks for the disks) on my system. This superseded MDK 8, the install went OK, but although set up for the KDE desktop and the ISP details all entered correctly, *kppp* was not installed. Trying to install *kppp* from the same disks with the software manager didn't work, it produced the following error:

```
kdenetwork-kppp-3.1.3.37
mkd.i586.rpm: invalid
signature((SHA1)DSA sha1 md5
(GPG) (MISSING KEY) GPG#
70771ff3 NOT OK)
```

Do you want to continue?

Presented with this wonderful set of hieroglyphics, I considered it best not to continue. I also looked for *Bluefish*, but this was not installed either, also receiving the same error when trying to install it. So is it safe to continue with the install? And if not, any suggestions as to what to do next? (Any Web access is 28KBps narrow band, thanks to BT...)

Also the system is set to automatically log on as a user (system is not networked and I am the only user); however, when

Most printers are compatible with your flavour of Linux – it's just a matter of finding the right drivers.



logging out from the KDE desktop there is no longer a shut down option, being returned instead to a BASH prompt, and I have to su to root and shut down from the command line. (Not very clever.)

How can I get the system to shut down direct from the desktop, or via a few clicks as with MDK 8?

Dave Pritchard, via email

A It looks like you're missing the public PGP key for whoever signed that RPM. You can force the RPM through and hope for the best, although it's possible that the RPM has been modified and the key, if verified, is now incorrect. You might want to look at ensuring your RPM package is up to date, as this should provide the GPG key required to verify the packages.

Normally, a non-root user is not able to shutdown the system, as it requires a change of runlevel. You can get around this by doing as root:

```
# chmod +s `which shutdown`
```

This will allow a non-root user to execute **shutdown -h now** to shut the system down and you can create a button in KDE to run this.

Old machine scan

Q I've got an old Pentium with a Slackware install (no X11) and I want to change it into a router. Also I'm thinking of buying a HP PSC1210 printer+scanner and while I'm at it I'd like to make the same Pentium also a printer and scanner server, so I can print and scan from all of the machines I have

at home. The question then is how (and if) this would be possible? I'm especially worried about the scanning bit, I know there's enough information out there for the routing and printing.

Matija Suklje, via email

A Printing is easily done, as both *CUPS* and *lpd* support remote printing. You can also set up printing for Windows systems using *Samba*, which essentially makes the Linux system provide Windows shares to the network. Information on this can be found in the SMB HOWTO at <http://tdp.org/HOWTO/SMB-HOWTO-9.html>. You can also find out all about Linux printing at <http://tdp.org/HOWTO/Printing-Usage-HOWTO.html>. Depending if you've got *CUPS* or *lpd* installed already, the first step is always to get local printing working, then work from there with remote systems.

Scanners generally interact via the SANE system (www.sane-project.org/), although the HP device you mention is currently not supported. SANE does support network scanning, which is rather useful, and docs on this can be found at www.xs4all.nl/~ijm/SANE-faq.html#92.

There is a great list of supported devices on the SANE site, so you would be advised to work through that list prior to making any scanner/printer purchases, to ensure that you don't end up with something that does not work with Linux. **LXF**

missed one?

LINUX FORMAT BACK ISSUES

Every month *Linux Format* brings you the best tutorials, the essential reviews and the latest news. But if you've missed out on a must-read feature or a vital program from our expertly compiled CDs and DVDs, order your back issue NOW! And remember, you need never miss an issue of your favourite Linux mag, if you subscribe to *Linux Format* (see overleaf for more details).



Christmas 2003

Product code:
LXFB0048(cd)
LXFD0048(dvd)

DVD HIGHLIGHTS:
 KDE3.2 alpha2,
 DamnSmallLinux,
 Audacity, Samba 3.0

CDs HIGHLIGHTS:
 SUSE Live Eval 9.0
 (boots from CD/ISO),
 Slackware 9.1, gentoo
 File Manager, Karchiver,
 BitTorrent, Mozilla,
 Opera, Firebird, Mambo
 Open Source Site Server

MAGAZINE FEATURING:
 OpenOffice.org & StarOffice,
 HP Itanium server, KDE3.2,
 MIDI sequencers, BitTorrent,
 Majesty review, Scribus DTP
 tutorial, Slackware 9.1 review,
 Gigabyte SR125ED review



December 2003

Product code:
LXFB0047(cd)
LXFD0047(dvd)

DVD HIGHLIGHTS:
 OpenOffice.org (PPC),
 MailScanner, XFCE,
 LDP, PuppyLinux,
 Slurm, fiphoto,
 Battlemech, XMMS,
 Fandango, CVSToys,
 Scrubber, Sylpheed

CDs HIGHLIGHTS:
 Mandrake 9.2,
 OpenOffice.org,
 GNOME 2.4

MAGAZINE FEATURING:
 Mandrake 9.2, Linux Expo UK
 report, Lazarus, CVS, Xinit
 SharpStation, MTStudio, SUSE
 9.0 review, Scanners tutorial,
 Newsreaders roundup, IBM's
 Linux strategy, Nessus



November 2003

Product code:
LXFB0046(cd)
LXFD0046(dvd)

DVD HIGHLIGHTS:
 RUNIT, 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:
 Topologilinux, 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

Other back issues also available, please call for further information, giving the issue number and/or date required

To order, call the back issues hotline on **0870 4448645**

or **+44 1458 271178** for overseas orders

Remember to quote the issue code shown above and
 have your credit card/switch card details ready.

Prices (inc post and packing): CD – UK£5.99,
 US & Europe £6.99, Rest of world £7.99
 DVD – UK £6.99, US & Europe £7.99, Rest of world £8.99.

LINUX FORMAT

www.linuxformat.co.uk



SPECIAL SUBSCRIPTION OFFER!

SAVE 40%* ON THE PERFECT GIFT!

A subscription to *Linux Format* makes a fantastic present for a friend, relative or even yourself!

- ✓ Save over **£31*** on the price you'd pay in the shops
- ✓ **FREE**** gift card for you to give
- ✓ **FREE** delivery direct to your door
- ✓ **Never miss an issue!**

LINUX FORMAT BRINGS YOU:

- 1** Essential reading for all Linux users
- 2** 20-page *Linux Pro* supplement with every issue
- 3** Three CDs or a DVD packed with Linux software
- 4** Latest news, latest reviews and the best Linux tutorials around

Taking out a subscription or giving a gift is easy. Simply call **0870 444 8645** or complete the order form opposite and send to the freepost address.

*40% Saving is based on 13 issues bought on the UK newsstand, this offer is only available to delivery addresses within the UK. £31 saving based on price of DVD issue.



ORDER HOTLINE ☎ 0870 444 8645

User Groups

LUGs worldwide are full of members keen to help with your problems, discuss ideas, and generally natter about all things Linux. You can find lots more information online at: www.lug.org.uk

1 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.malvernlug.cstf.co.uk
Contact Phil Ironside – phil@cstf.co.uk

REVISED
DETAILS

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



LINUX USER GROUPS

**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

LINUX USER GROUPS

LUG OF THE MONTH

Schoolnet Namibia

Matt Nailon writes: A year ago, this organization providing much-needed computing resources turned down Microsoft's offer to put Windows systems in several of its schools and decided to continue using Linux systems. Rather than let Schoolnet.na become a PR tool for Microsoft, Executive Director Joris Komenn publicly lambasted Microsoft for a plan that would give the schools a \$2,000 break on *MS Office* software but make them pay US\$9,000 for Windows XP.

Schoolnet.na has evolved into a crucial advocate and example of how Free/Open Source software (FOSS) schemes can be very successfully implemented in a variety of educational environments. Regarding the difficulties that teachers in the UK and US are experiencing in convincing decision-makers that FOSS is the way forward, www.schoolnet.na shows that Westerners are wrong to describe Namibia as being a 'developing country' – obviously, we are the nations that are in dire need of catching up.



Worldwide Linux User Groups

Free Software users across the globe

Africa

EGYPT

URL www.linux-egypt.org

GAUTENG, SOUTH AFRICA

URL www.glug.org.za

Email glugmin@revolution.org.za

THE LORD'S ABODE, JO'BURG, SA

Email Andrew Gargan avrin17@iname.com

Australia

ADELAIDE

URL www.linuxsa.org.au

Email mtippet@anu.edu.au

ALICE SPRINGS

URL www.aslug.org.au

MELBOURNE, VICTORIA

URL www.luvasn.au

Contact luv-committee@luvasn.au

PERTH

URL <http://plug.linux.org.au/>

SYDNEY

URL www.slug.org.au

Europe

COSTA DEL SOL (English speaking)

URL www.fuengirola.lug.org.uk

DENMARK

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.ed

Asia

HONG KONG (multilingual)

URL www.linux.org.hk

SINGAPORE – SLUG

URL www.lugs.org.sg

SRI LANKA

URL www.lklug.pdn.ac.lk

MYANMAR (formerly BURMA)

URL www.myanmarlug.org

Email afyde@balug.org

PAKISTAN

URL www.linuxpakistan.net

Email tux@clug.org

HYDERABAD, SINDH, INDUS VALLEY

URL www.geocities.com/slug_pk/

KASHMIR

Coming soon!

China

BEIJING (GB encoding, but mostly written in Chinese)

URL <http://mud.263.net.cn/~linux>

CHINESE LINUX USER GROUP

URL www.linux.org.cn

NANJING

URL <http://jllib.jlonline.com/njlug>

India

LINUX INDIA

URL <http://linux-india.org>

ALIGARH LUG

URL <http://linux.amupost.com>

BOMBAY

URL www.ilug-bom.org.in

CHANDIGARH

URL www.geocities.com/vipinb

CHENNAI AND MADRAS

URL www.chennaiug.org/

CYBERABAD (CLUG)

URL <http://seeknew.freesevers.com/clug/>

DELHI

URL www.linux-delhi.org

KOLKATA

URL www.ilug-cal.org

MADURI

URL <http://linuxmadurai.tripod.com>

NORTHERN INDIA LINUX

URL <http://groups.yahoo.com/group/lug-northindia>

Spreading the word

“How can Linux make *my* business more profit?” Jono Bacon explores the pitfalls of advocating Linux use in enterprise.

One critically important issue connected to Total Cost of Ownership (TCO) is staff retraining and availability. Although Linux has proven successful as a server-based OS, the issue of training can be a thorny one, particularly when applied to legions of desktop users in a larger organisation. Even for sysadmins, operating and maintaining Linux is considerably different to Windows systems. Traditionally, Unix users have had a much easier time migrating as there is less to learn in terms of system administration.

When advocating the training issue for technical staff, you need to realistically look into the Linux talent in your area for help, and even the employment of specialists for training and advice. One thing a business does not want is to be left in the lurch with an OS that no one can administer. There are bound to be experienced Linux admins locally, and LUGs are a good source of contact. Also look on the normal job boards for Linux admin CVs. There are many companies that can train people to administer Linux systems, including Red Hat. Cost for this

kind of retraining is mostly comparative to that of Windows-based training, and the savings with a Linux solution could easily be seen to write this cost off.

The question of desktop user training has to be examined on a job-description basis. If your company is deciding that that the *OpenOffice.org* suite would make an appropriate workable replacement for *MS Office*, staff will need to be trained in not only how to operate a replacement program itself, but also how it interacts with proprietary formats, and what shortcomings may be encountered in tasks they will regularly be performing in their role. This goes against the current business IT use model somewhat, where many companies assume that the bulk of administrative and accounting staff will bring computer skills to the job. The extra cost factor here needs to be offset by demonstrable gains in productivity and efficiency brought about by users having a better understanding of the systems.

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.

CREATING INSTALL CDS WITH CDRECORD

The quickest way to burn an ISO image to CD is with *cdrecord*. You need to be root to do this. First find the address of your CD-writer with

```
cdrecord -scanbus
```

This will show the devices connected to your system. The SCSI address of each device is the three numbers in the leftmost column, say 0,3,0. Now you can burn a CD with

```
cdrecord dev=0,3,0 -v /path/to/image.iso
```

You can simplify the command by saving some default settings in `/etc/default/cdrecord`. Add a line for each CD writer on your system (usually one) like this

```
Plextor= 0,3,0 12 16M
```

The first item is a label, after the SCSI address you put the speed and the buffer size to use. You can now replace the SCSI address in the command line with the label, but it gets even easier if you add

```
CDR_DEVICE=Plextor
```

Now you can burn an ISO image to disc with

```
cdrecord -v/path/to/image.iso
```

If you really don't want to use the command line, *gcombust* will do the job for you. Start it as root, select the "Burn" tab and the "ISO 9660 Image" gadget near the top of the window. Put the path to the image file in the gadget and press "Combust!". Now put on the kettle while the CD is created for you.

Other OS?

You don't have to use Linux to burn the ISO to a disc. All Linux-specific bits are already built into the image file. Programs like *cdrecord* simply dump it to the disk. If you don't have a CD-writer, find someone who has one, and a DVD drive, and use the CD burning software on their computer. It can be Windows, MacOS, AmigaOS whatever.

No CD burner?

What if you have no CD writer? Do you know someone else with one? You don't have to use Linux to burn the CDs, any operating system that can run a CD-writer will do the job (see above).

With some distributions it is also possible to mount the images and do a network install, or even a local install from another disk partition. The methods often vary between distributions, so check on the distro vendors website for more information.

WHAT ARE ALL THESE FILES?

If you are new to Linux, you may find the profusion of different files and extensions confusing. As we try to give as many packages as possible for compatibility, there will often be two or three files in a directory covering different types of Linux, different architectures and usually source and binary versions – so which do you install? They can be identified by their filenames, and usually just by the file extensions.

Someap-1.0.1.i386.rpm – This is probably a binary rpm, designed to run on x86 systems.

Someap-1.0.1.i386.deb – The same, but a debian package.

Someap-1.0.1.tar.gz – This is usually source code.

Someap-1.0.1.tgz – Same as the above, tgz is abbreviated form of targz

Someap-1.0.1.tar.bz2 – Same, but uses bzip2 compression instead of zip

Someap-1.0.1.src.rpm – This is also source code, but supplied as an rpm to make it easier to install

Someap-1.0.1.i386.RH7.RPM – A binary, x86 RPM designed specifically for Red Hat Linux

Someap-1.0.1.ppc.Suse7.rpm – A binary RPM designed specifically for SuSE7x PPC Linux.

Someap-devel-1.0.1.i386.rpm – A development version.

INSTALLING FROM TARBALLS

A tar ball is a two stage archive. First the files are archived into a single file with *tar* and then compressed with *Gzip* or *Bzip2*. To unpack, *cd* to the directory you want to unpack it, usually your home directory and type one of the following two lines:

```
tar xzvf /mnt/cdrom/Desktop/progname/progname-2.1.0.tgz
tar xvf --bzip2 /mnt/cdrom/Desktop/progname/progname-2.1.0.tar.bz2
```

Use the first for Gzipped files, those ending in `.tar.gz` or `.tgz`, and the second for Bzipped files, ending in `.tar.bz2` or `.tbz2`. Naturally, you change the paths to suit the location and name of the archive. and replace `/mnt/cdrom` with whatever is applicable to your system (eg `/cdrom`). This normally unpacks the archive into a directory of the same name, enter that directory with:

```
cd progname-2.1.0
```

To compile and install the software, type the following three commands:

```
./configure
make
su -c "make install"
```

The last line will prompt you for the root password, as this stage must be run as root. If you are already logged in as root, just type **make install**. This will give you a default installation. If you want to change any aspect of the install, type **./configure --help** to see the options available. For example, you are usually able to change the default location with the PREFIX argument. When you have finished installing, you may remove the source files with:

```
cd ..
rm -fr progname-2.1.0
```

You should also log out as root, before you do anything you may later regret.

DEFECTIVE CDS

In the unlikely event of your disc being defective please email our support team (support@futurenet.co.uk) for further assistance. If you would prefer to talk to a member of our reader support team please call **01225 822 743**.

COVERDISC DVD

« select a file from whichever directory you configure and oooqs maintains a list of your most recently opened documents, adding them to the menu.

One of the most useful features of oooqs is the option to preload and

hide a copy of *OpenOffice.org* when it starts up. Obviously, your computer will need plenty of spare memory for this to work effectively, but it then means that the program opens almost instantly when you need it.

DISTROS MANDRAKE

When we included Mandrake 9.2 on the cover discs, there was no kernel-source package, because we were not supplied with it. This package is

important, as you usually need it when compiling software from source. Its omission won't prevent you installing any RPM files, either from a Mandrake repository or generic packages, but there are programs that are only

DVD CONTENTS AT A GLANCE

Desktop

ACPIClient
BrutalFileManager
CDDetectExecute
DragStack
EmelFM2
GPA
Informant
K3b
KickPIM
KNetworkConf
LDAPAccountManager
LDAPExplorerTool
Luma
MFEEditor
NEdit
Ocrad
Regexer
ShutDown-O-Matik
SMBfax
TVInABox
Wget

A command line tool that provides useful ACPI information
A 3D FPS filemanager
A CDROM disc type detector
A kicker applet on which you can drag files
GTK+2 file manager that implements the two-pane design
GTK frontend to GnuPG
A GNOME panel applet useful for notifying a user of events
CD-Burning software built using KDE
A KDE Kicker address book applet
Configure TCP/IP settings under KDE
A PHP4-based account manager for openLDAP
An LDAP browser
A utility for accessing and managing LDAP data
A simple utility for renaming multiple files simultaneously
Advanced multi-language programmers' text editor
The GNU OCR program
A nifty GUI search/replace tool using Perl regexps
Shut Down-O-Matik is a graphical shut down utility for KDE
Sends a fax to a Samba print share without Win32 software
An applet that shows TV
A network utility to retrieve files from the Web

Development

ACPI.py
ClanLib
CommandLineProgressBar
FOX
Hermes
ImageroReader
IzSound
Iidcheck
libDSP
libstatgrab
PerlArchiveToolkit
Pygame
Qt
VarnasAutomotiveSimulator

A Python module covering ACPI
The platform-independent game SDK
Display information about a data transfer stream
A C++-based library for graphical user interface development
A CRM and business tools development environment
A Java library for reading images
A C++/STL sound processing library
A Linux dependency checker
A C++ library of DSP routines
A useful interface to system statistics
Tools for Perl packaging and deployment
A Python game development package, based on SDL
A GUI software toolkit
An automotive simulation framework and application

Distros

LocalAreaSecurityKnoppix
Warewulf

Live CD distro focused on security and small footprint
A distributed Linux distribution

Games

HuntForGold
Oddspace
QuantumStarSE
SDLJoytest
Seriken
TuxMathScrabble
Ultrasol

A real-time strategy game in Java
A multiplayer 2D space shooter
Space strategy browser-based game
An SDL-based joystick testing program for X11
A city simulation game
A math scrabble game for all ages
A collection of solitaire/patience games

Graphics

DVDrip
FFTV
Frameworks
Inkscape
Mkxvcd.sh
Motion
Ogmitools
PP3
Transcode
Wbcopy
Wings3D

A full featured DVD Ripper GUI
A viewer and recorder for TV and radio
Webcam software for stop-motion animation
An SVG-based vector drawing application
A script to create a 2 hour VCD
A motion detector for video4linux devices
Tools to merge different streams into Ogg files
A tool to generate celestial maps with LaTeX
A video stream processing tool
Copies DVD .vob files to harddisk
A polygon mesh modeler

Help

LDP

A complete mirror of the Linux Documentation Project

Internet

AmritaVPN
Anamail
Brag
ChestnutDialer
Copernicus
EchlotPinger

An open source SSL-based Virtual Private Networking tool
An email statistics collection and analysis system
Download and assemble multipart binaries from newsgroups
A PPP dialing program
A program for remembering usernames and passwords
A pinger for anonymous remailers

Gaim
Gwget2
Hastymail
Kopete
MPlayerplug-in
Pink
Pushmail
SvenDeleter
TigerEnvelopes
Tinc Virtual private network daemon

Mobile

AccessPointUtilities
Wifimap
WmWiFi

Configure and monitor a Wireless Access Point
Document the structure of a wireless network
A wireless network signal strength monitor for WindowMaker

Office

AbiWord
Avsap
Brinace
CybercafeOrganizer
DocumentManager
KSqlAnalyzer
Scribus

A fully-featured word processor
Powerful, flexible and complete financial accounting
A command line financial planner
A cybercafe point of sales and administration tool
A multi-user document management system
Work on an MS SQL database server
A DTP-Program for Linux using the Qt-Library

Server

Anteater
Apache
Dnsmasq
ErlurtWiki
Httaccesser
htDig
Html2Wml
IMAPProxy
InsideSystemsMail
mod_ssl
MySQL
OneTimeDownload
Open3SP
Pathalizer
PieterPost
Renattach
Scrubber
ScyllaCharybdis
WTP
YetAnotherAntiVirusRecipe

A modular sendmail and postfix mail log analyzer
A high performance Unix-based HTTP server
A simple lightweight caching DNS forwarder
A fast, user-friendly and highly configurable Wiki library
A script for generating Apache .htaccess files
Complete world wide web indexing and searching system
A program that converts HTML to WML
A caching IMAP proxy server
A web-mail system for IMAP accounts
Apache Interface to OpenSSL
A fast SQL database server
A script to allow downloading files using a one time ticket
Secure application framework server providing VPN and IM
Visualise the paths users take when browsing a Web site
A Web-based PHP frontend for POP3 mailboxes
A filter that renames or deletes dangerous email attachments
Server-side mail content filter system
Compressed, encrypted, throttled file upload/backup
A Web-based FTP client
Procmal recipe to filter out most common e-mail worms

Sound

ALSA
Audacity
ComOnFeelTheNoize
cpShuffleSature
Gnomadic
GStreamer
Moosic
MPFC
MusicManager
Qjackctl
UnattendedBroadcastDaemon
XMMS-KDE

An alternative implementation of Linux sound support
A cross-platform multitrack audio editor
Generates an M3U list from any Ogg/MP3 on a network
Copies random files (MP3, OGG, MPG... anything)
GUI to control a Creative Nomad II/c MP3 player
A streaming media framework
A powerful music queue manager
A music player for the Linux console
A Konqueror plugin for management of audio files
Control the JACK sound server daemon
A daemon for running the operations of a radio station
MP3 player and a XMMS and Noatun frontend for the KDE panel

System

AddUser-NG
BitDefender
Captive
ClamAntiVirus
CryptoFS
Ethereal
Gimp-print
Gmodconfig
Nessus
NetworkComputers
Rsnaphot
Usermin
Webmin

Adduser script with more functionality, such as plugins
A command line antivirus scanner for Linux
Full read/write access to NTFS disk drives
An anti-virus utility for Unix
An encryption filesystem for LUFES
A GUI network protocol analyzer
Top quality printer drivers for POSIX systems
A tool for manipulating kernel modules
A free, open-sourced and easy-to-use security auditing tool
Inventory software for network administrators
A filesystem snapshot utility
A Web-based interface for UNIX users
A Web-based interface for Unix system administration

supplied as source. Mandrake has updated the kernel package since the release, so we have the new versions of both the binary and source packages on the cover discs this month.

Installation of the kernel itself is slightly different from other packages, in a root terminal type


```
rpm -ihv /mnt/cdrom/Distros/
Mandrake/kernel-2.4.22.21mdk-1-
1mdk.i586.rpm
```

The kernel source should be installed with *urpmi*, as it does have some dependencies. *urpmi* will prompt you for **LXF DVD47** and install these for you, so type

```
urpmi /mnt/cdrom/lxf/dvd/Distros/
Mandrake/kernel-source-2.4.22-
21mdk.i586.rpm
```

INTERNET KOPETE

Instant messaging (IM) is one of the largest growth areas of the Internet right now. There are even some pundits claiming that it will become a greater killer application than email. If you tried out the KDE 3.2 alpha version from last month's DVD, you



Instant Messaging KDE-style. Kopete will be an integral part of KDE 3.2. Here it is for those of you who would rather not wait.

may have already seen *Kopete*, as it is part of KDE 3.2. *Kopete* is a universal IM client. It uses a plug-in system to deal with the various IM protocols. This

CREATING INSTALLATION CDS FROM THE DVD

As is now normal with our bootable DVDs, we have provided a way to create CD ISO images so that you may burn your own CDs for installation on a computer without a DVD drive. This can be done from Linux or Windows. We have made some additions to the Linux script, as a result of feedback from users on our web forum.

To build the ISO images in Linux, type the following command in a terminal

```
sh /mnt/cdrom/Distros/Fedora/mkiso
```

This will create the three ISO images in the current directory. If you want to create them somewhere else, give the path as an argument, for example:

```
sh /mnt/cdrom/Distros/Fedora/mkiso
/tmp/iso
```

Note that you should not **cd** to the directory on the DVD when running this script. It will not fail as it used to, but it will be slower because the script is not able to open the cache file used to speed up the creation of the second and third discs.

If you are short of space, it is now possible to create single ISO images with either:


```
sh /mnt/cdrom/Distros/Fedora/mkiso -d 1
or
sh /mnt/cdrom/Distros/Fedora/mkiso -d
2 /tmp/iso
```

Windows users can create the CDs by double-clicking the *winkiso* icon or alternatively running the script from a MSDOS prompt.

In the latter case, you should change to the Fedora directory of the CD before running the script. The Windows script does not allow you to create single ISOs, but you can specify a destination directory when running it from a DOS prompt. With no argument, or when run from the icon, it puts the ISO images in C:

For example, if your DVD drive is designated 'E:' and you want to save the ISO images to D:\ISO

```
E:
cd Distros\Fedora
winkiso D:\ISO
```



After creating ISO images with the *mkiso* script, you can burn them to CD with *K3b*, which just happens to be on this month's DVD.

DVDS: TO BOOT OR NOT TO BOOT

Smart Boot Manager



Use Smart Boot Manager to persuade a reluctant computer to boot from your Linux Format DVD.

Although many of our DVDs – including this month's and the recent Mandrake and Slackware DVDs – are made to be bootable, some people do seem to have problems booting from them. Here is some information that has been posted on the *Linux Format* web forum about the boot process.

*“There are two common ways of making a CD or DVD bootable. The original method was to embed an image of a bootable floppy disc, which in turn started the CD/DVD boot process. This has the advantage of working on almost anything, but limits what you can boot to the size of a floppy disk. The later method uses *isolinux* instead of emulation. It offers more flexibility to the discs' creators, but won't boot on some hardware. You can tell if a disc uses *isolinux* by the presence of an *isolinux* directory, usually (but not always) in the root of the disc.*

“The choice of booting system is not made by LXF but by the creators of the original CDs on which the DVD is based.

“DVDs boot in exactly the same way as CDs, they are treated exactly the same, they just hold rather more. However, I have seen occasions when a DVD will not boot while a CD using exactly the same boot files will. I suspect that this is due to the longer time taken to

recognise and handle a DVD by the drive, so the BIOS has already decided the disc isn't even there before the bootable information can be accessed. In this case, delaying the boot process long enough for the drive light to stop flashing, by going into the BIOS menu and then exiting without changes, can make the disc boot.

“In other words, it is not that the disc is not bootable, only that your computer will not boot from it. I agree that creating a boot disk for every DVD is a pain, and some require a number of discs, which is why Smart Boot Manager is in the Essentials directory of the cover discs. Make a boot floppy from this and you should be able to boot any of the DVDs.”

One flaw with this suggestion is that some computers have no floppy drive; so we have made a CD ISO image, called *shootmgr.iso*. Burn this to a CD-R and boot from that. When you see the boot menu, eject the CD and replace it with the DVD before selecting the relevant device (usually CDO). We have also made a minor change to the way in which the ISO images are created that may make a difference. If you could not boot the Mandrake or Slackware DVDs but this month's Fedora disc does work, please let us know.

means new system can be quickly and (usually) easily added, and it doesn't even have to be done by the *Kopete* authors. Anyone could release a new plugin as the API and even some plugin templates are included with the package – isn't the Open Source way of doing things wonderful?

Incidentally, some people were unable to compile and install the KDE 3.2 alpha from last month's disc because they did not have access to a recent enough Qt package. So this issue we have included the very latest version, 3.2.3, in the Development directory of the DVD. [LXF](#)

LINUX PRO

FROM THE MAKERS OF LINUX FORMAT

JANUARY 2004



PLUS

Open Minds

100 per cent availability may not be a pipe dream if you're talking to the right people

Filesystem Virtualisation

Making your storage solution seamless

Shado content management

ColdFusion MX and Linux combine to deliver for London Metropolitan University

Hosting profile

Hostway Europe's Linux services in the spotlight

The BLADE advantage

Cost saving, consolidation, high-density computing, cable reduction, modular, versatile, high availability – is it time to reassess blades?

PRACTICAL LINUX SOLUTIONS FOR I.T. PROFESSIONALS

Welcome

Twenty pages of real-world Linux for IT professionals

There seems to be a slight undercurrent of an 'availability' theme to this issue, which isn't too surprising. Whereas in the past people expected the odd 'computer failure' and the Internet seemed to be 'broken' more often than not, attitudes are changing. If you can't keep your servers up, why should people trust you with their business?

Deciding how far you need to go with availability processes can be a tricky business. There's no point, for instance, in spending large sums on doubling up servers if they share the same data source. Not only do you have to factor in the likelihood of various failures, but also how long they would take to rectify. Unsurprisingly, many people turn to specialists in these matters, and we take a look at probably the best-known of these organisations Open Minds on page 12.

If you are seriously investigating high availability (HA), you may want to take a quick gander at the eServer BladeCenter that we showcase over the page. Blade technology has certainly matured now that their advantages are more clearly understood. You can see how the very latest IBM HS20 blades shape up over the page.

As ever, there is plenty more in this issue – we investigate the different options for filesystem virtualisation in our regular storage section, and there is another profile of a web host. This issue we are examining the services offered by Hostway, and chat to some of its customers. And don't forget to check out the short case study illustrating how London Met University managed to meet a tight web deadline with the help of Linux and a new Content Management system from down under. Enjoy!

Nick Veitch Editor
nick.veitch@futurenet.co.uk



“Not only do you have to factor in the likelihood of failures, but also how long the problem will take to rectify – many people turn to specialists...”

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

Contributor

David Cartwright

Publisher

Kelley Corten

Group Publisher

Dave Taylor

Photography

Neil Godwin, Photodisc, Corbis

ADVERTISING

Senior Sales Executive

Andrew Hudson

andrewhudson@futurenet.co.uk
(01225) 442244 Ext 2441

Copyright No part of this publication may be reproduced without written permission from our publisher. All letters & contributions to *Linux Format* are submitted and accepted on the basis of non-exclusive worldwide licence to publish or license others to do so unless otherwise agreed in advance in writing. *Linux Format* recognises all copyrights in this issue. Where possible, we have acknowledged the copyright holder. Contact us if we haven't credited your copyright and we will always correct any oversight. We cannot be held responsible for mistakes or misprints. **LINUX** is a trademark of Linus Torvalds, GNU/Linux is abbreviated to Linux throughout for brevity. All other trademarks are the property of their respective owners.

LinuxPro is a part of



magazine

Published by



Media with passion

30 Monmouth Street, Bath BA1 2BW
Tel +44(0) 1225 442244

JANUARY 2004

Contents

BLADES

The latest Blade solution from IBM on test **p04**

LONDON MET

Implementing CM with *Shado* **p06**

VIRTUAL FILESYSTEMS

Storage strategies explained **p08**

OPEN MINDS

Maintaining high availability with the experts **p12**

HOSTING

Focus on Hostway **p18**



4

12

LINUXPRO 3

Blades of steel:

Is this the future of computing?

Like Grid Computing, the blade server was an industry-redefining solution that is starting to make its presence felt more than ever. For the Grid paradigm, computers across distributed networks are grouped together to tackle a complex task faster and better than the machines would be able to do alone. While the Grid way of thought certainly has its uses, and is indeed likely to become ubiquitous in the future, it's much too large a solution in its current state to allow companies to use it for their own needs. Companies need ultra high-performance computing right now to handle tasks such as web serving, file sharing, or pretty much *any* task where super-dense clustering is required.

Poor man's supercomputer?

An average blade rack will hold between 10 and 16 blade servers, and each blade server will usually be either uniprocessor or dual-processor, which means a full blade rack will generally have between 10 and 32 CPUs on-board. It's common knowledge that the only way to get a 6.0GHz Intel computer today is to buy a dual 3.0GHz box, but this principle only scales up neatly so far as the price escalates exponentially. Past four-way machines, relying on symmetric multiprocessing (SMP) machines becomes less economical, which meant another alternative was required.

This is where blade comes in – hundreds or thousands of low-price machines can provide a great deal more computing power than would otherwise be possible using SMP. Google uses this principle to provide its gigantic search capabilities without requiring huge amounts of dollar bills to back it up, as each box costs a fairly small sum while making up part of a shared machine that is able to outpace supercomputers.

Because blades are just normal servers optimised to take up as little space and power as possible, they can theoretically be put to any task you want. Of course, some might say it's a waste to have a 72-machine blade rack running as a file and print server, but it's certainly possible. The key to eliminating much of the space, heat, and power is in the configuration of each server – they often use fan-less CPUs, are fitted with 2.5-inch hard disks, and have their RAM sticks laid almost horizontally, which all contributes to the high-density nature of the overall system.

Many blades are actually hot-pluggable, which means that you can juggle your servers around as necessary without any downtime. Want to increase your render farm by 10 per cent at the expense of your web server? Simply take out blades from the web server group as necessary and plug them into the render group – it's that easy. This easy divisibility means that blades have quite a substantial benefit other than cost when they are compared against traditional supercomputers.

High-performance computing is driving the industry in a new direction, but do blade servers really fulfill the need they were designed for, asks PAUL HUDSON...

This divisibility can also, paradoxically, help unify your system administration. For example, it's easy to set up one set of servers to act as your live system and another set, in the same rack unit, to act as the backup system, thus encapsulating your failure recovery plan in one place.

Day walkers

Blade servers themselves are fairly simple devices, not much removed from standard servers; however, when they plug into the blade rack device they gain a higher level of intelligence because they usually share a common bus at the back for intercommunication purposes. Saying "the whole is greater than the sum of its parts" is quite fitting when describing how the blade racks work, as they provide the power and networking controls to the blades to complete the equation and make the system as powerful as it is. As you can imagine, the blade racks are where the real complexity comes in from a technology perspective, and is where the largest chunk of the cost comes from.

Big complexity usually means big name brands powering the systems, and so it's not surprising to find IBM, Sun, HP, Fujitsu-Siemens, Intel, Dell, and other big-name brands all working overtime to promote their blade offering over everyone else's. The competition is hot as the market is surging forward, which means you're likely to get better prices if you shop around.

To give you an example of what to look for in a blade server, we asked IBM to send us over its renowned eServer BladeCenters for review, and it arrived just a few days later. At 7U and supporting up to 14 blades per BladeCenter, this was certainly no small piece of hardware. Each of the blade

"Hundreds or thousands of low-price machines can provide more computing power than would be possible using SMP!"

servers had a 2.8GHz Xeon CPU, a 40GB hard disk, and 512MB RAM as standard, which brings them in at the low price of £1,390 each. The BladeCenter unit itself costs £2,015, so as it came with four blades the total price was £7,575. To extend this to the full capacity of the server, the total cost would be £21,475 before VAT, giving you a total of 40GHz of computing power for just over £20,000. Of course, this kind of value would be impossible if each server had to run on a non-free OS!

Although the price is likely to be negotiable, particularly if you do opt to buy several blade servers or even several BladeCenters, you're still getting an awful lot for your



money. Consider the HP Integrity rx2600 that we reviewed last issue – it came in at £22,930+VAT with just 3GHz powering it. Now, no matter how much faith you have in the Itanium chips used in the rx2600, it's just not likely that they will perform ten times faster than equivalently clocked Xeon chips, and, *even if somehow they did*, that would still only make the Itanium equal to 30GHz of Xeon computing power, and therefore *still 25% slower* than the BladeCenter.

Technical Updates

Opening the BladeCenter packaging, you can't help but concentrate on the clever way IBM have made the blades lock into the main server, using a dual-catch lever system to make sure the servers are firmly in at all times and yet only requiring a light pressure on both catches to release them.

Once the main server is taken out of the box, the first thing we noticed was the pile (yes, the *pile*) of manuals. Sifting through them we noticed the large majority were marked 'Technical Update', which included a description of fixes produced by IBM since the product was launched. This might seem a little worrying at first, but the majority of the fixes were applied for us already, with others coming with clear instructions. If the patches were applied already, you may well wonder where the point was in telling customers about it, but it actually works out for the best – you can tell at a glance as soon as you open the box precisely in what state your servers are in, and what you need to do to maximise your security.

This easy level of maintenance is backed up strongly by IBM's corporate warranty – 3-year on-site parts and labour as standard – but you can always arrange with IBM to extend this if necessary. This comes back to one of the key advantages to blades – if any of the servers fail, you can just swap it out and carry on working regardless while the

It may look like a Borg Cube™, but there's not even a whiff of Microsoft anywhere inside, and will crunch your numbers at speeds approaching Warp 9.

original blade is being serviced. To get this same kind of failover capacity with a traditional supercomputer is very difficult, and will certainly leave a bigger hole in your wallet.

Blade runners

So, with a massive amount of processing power taking up a fairly slim spot in a rack at what work out to be a very cost-effective price, are blades the most power for your money? Well, no – you can get white-box servers for cheaper than this if space is of no concern, but this is rarely the case. The extra price you pay for blades is caused by the high-density factor, which is often critical when you're working with large numbers of machines.

To render *Lord of the Rings: Return Of The King*, Weta Digital needed 1600 dual-CPU servers to get everything done in time. In the high density world of blade servers, this would still require 115 fully loaded BladeCenters, which would certainly take up a lot of room, but consider how much space that would take up if it *weren't* running on blades!

This means that blades solve a key problem in computing right now, which is that the need for more and more power is pushing the size of computers above and beyond acceptable limits in environments where space is limited. By keeping space and power requirements to a minimum, blades are easily stacked up to produce systems of a power that would have been impossible just a few years ago at a price that would have been impossible just a few *months* ago.

Blade servers, particularly those powered by Linux, seem to be surging forward into every market right now, and we hope to bring you more news and reviews over the coming months. With some competing operating systems costing almost as much as a blade itself, it seems that Linux has found itself yet another niche in which to excel... ■

CASE STUDY

Defending the skies with

content management

The name *Shado* may conjure up images of UFOs and Interceptors, it is actually the name given to New Zealand-based Straker's content management system. The world is not short of such software, which ranges, at least in the Linux market, from various Open Source efforts to cross-platform high-end commercial solutions.

To better gauge the effectiveness of the software in a real-world context, *Linux Pro* met up with Mike O'Reilly of London Metropolitan University, one of the first organisations to try the new version of *Shado*. Following a merger with another university, London Met needed to re-assess its web strategy. The situation presented a problem – and an opportunity – in terms of the online presence of the institutions. The two sites, and two very different methodologies would have to become one in a radical shake up to the way online content worked. All within an extremely tight deadline of ten weeks!

Mike O'Reilly explained: "Two major universities were going to be merged. There hasn't been a merger like that in British higher education. We were faced with two large universities with pre-existing websites, both of which were

If asked to merge two distinct online identities, many web developers would run and hide. NICK VEITCH finds London Met far from scared.

London Met University – online in time thanks to *Shado*.

based on Solaris running *Apache*, both were using CGI. Here we used mainly ColdFusion, the others were using PHP. We were then faced with the prospect of having to create a new website for the new university.

"Although we had two universities offering web content, the model by which that content was added was entirely different. We have always had a distributed content model, where we devolve down to departments and individuals. City campus used an external firm to design their main page, and had little or no tradition of individual departments adding content."

It all seems too familiar. A clash of cultures, sweeping new changes, unbelievable timescales. It soon became apparent to the London Met team that to continue with a template driven site would not be possible. It was clearly time for some more radical thinking.

"We decided that this was the moment we could innovate: we were going to go for a proper content management system. We have a very large knowledge management system here. We have been talking about how to manage content for a long time. There is a large amount of content in our virtual learning environment. We've always had the problem that stuck on the website is a bunch of HTML files that we couldn't manage.

"We want our systems to interoperate, we want to be able to impose standards, all sorts of things. Although we were short of time we decided this was the moment"

Problem-solving

Even the simplest content management system allows some form of workflow control and an easy to use interface to allow non-technical staff to add content. But which one to choose from the legion of options available?

"There are at least 100 Open Source ones. We looked at Perl-, PHP- and Python-based ones. We didn't stint ourselves: we actually looked at them *all*. Our instincts were that if there was a really decent Open Source one, we would consider that. We had a good look, particularly at the Java ones because of our skills in Java here. We decided that none of them would deliver out of the box."





Shado now uses ColdFusion MX technology to further enhance the feature-set.

website which needs to be running pretty much 24/7 we needed to go with a reliable operating system. Solaris is too expensive for something like that. Straker said that it was pretty committed to Linux, so we decided to put it on Red Hat, running on Dell servers. The Linux side was fine. Since we had a lot of Unix skills there was virtually no retraining. Because it's Linux it's been easy to implement"

The deadline met and the site up and running (check it out yourself at www.londonmet.ac.uk), the entire undertaking has been a success. But that's not the end of the story: "Our next challenge is FOI, or Freedom of Information. Any significant piece of information in this

It seems that the Open Source solutions, while being attractive and used by many, still don't have any great appeal for more corporate endeavours – yet. Part of this is down to the emphasis of the coding ("Open source tends to be really good at the front end, but terrible at the back end. When you start looking at how to add the users, the workflow – all the boring stuff – most of them would be impossible to manage on the scale we needed"). Still more is the lack of any support options for most of these solutions. Would you stake your job on being able to coerce such a system into working?

Once the field had been narrowed a bit, Mike O'Reilly's thoughts turned to the technologies on offer. Given the timescale, he decided it was prudent to stick with underlying technologies that the universities already had experience of using: "We came down to a clear choice between Java and ColdFusion. Given our experience with ColdFusion and the timescale involved, we decided to narrow it down to that. There were three or four contenders there, and we had some reps in to do a demo. We were attracted to Straker because it was the company most committed to ColdFusion MX. One thing we were convinced of was that Macromedia in the long-term will provide much of the multimedia capabilities we want right out of the box.

"Straker had existing Higher Education clients and they seemed to say the right things about collaboration"

Linux Pro also spoke with Tim Jenkins, Straker's UK MD, to find out about the commitment to Macromedia's new MX range: "The MX technology has put us in a good position. We were one of the first companies in the world to launch a product based on ColdFusion MX – it came out the week after MX was launched. We made that investment in time and effort when the economy was flat, and perhaps more by luck than judgement that seems to have been the right thing to do, as we are now positioned to take advantage of an upturn."

Choice of operating system

With the software sorted, there was still a choice of operating system to consider. Fortunately, according to Mike O'Reilly, this took little more than a few minutes thought: "Currently, we run Solaris. We were running ColdFusion on Win2000, and we were running various things on Linux. When we came to make the choice, well, the University is interested in Open Source. We pay a lot of money for Microsoft licences – we pay a lot for the server and desktops. All our experience of running all these boxes was that Solaris basically sits there and runs. Windows NT? You have to reboot them. Windows 2000 is a bit better. For the

"Since we had Unix skills, there was virtually no retraining. Because it's Linux, it's been easy to implement"

university that isn't related to you as a person – or has some commercial restriction – has to be made available. That has to really be made available through the website, so one of the issues we have is that we have to start flagging things with the metadata to indicate that it is FOI material. Then we'll need to build a new way to navigate through the site using this data. That's where we hope the benefits of this system will really pay off."

A framework

Tim Jenkins adds: "I think that's where we see the challenge for *Shado*. The way organisations are using data and the requirements they have for linking it up into one interface. We believe that *Shado* – as a framework – offers a lot of that facility now. One of the things we are keen to do is to get organisations talking to each other about how they can achieve these goals with *Shado*" ■

RESOURCES

www.londonmet.ac.uk/

Check out the site in action

www.straker.co.nz

More info on *Shado* and Straker's services.

The new site is just the start of the content London Met hope to deliver.



STORAGE

Implementing...

Disk Virtualisation

Disk virtualisation (or, to describe the process more correctly, *storage virtualisation*) is one of those concepts that seems to mean different things to different people. To some, a virtual disk is a RAID array; to others it's an NFS-mounted volume; to a third group it's a SAN-connected disk pack.

In fact, any storage implementation that puts some kind of abstract interface between a storage device and the system or person using it can be termed 'virtual'. The RAID array mentioned above has a number of disks that are made to look like a single volume via an abstraction layer. The NFS-mounted volume looks like just another directory on the local computer thanks to a different abstraction layer. Likewise, the SAN-connected disk pack, which looks just like a local disk to the computer's operating system, even if it's miles away in a data vault.


In this feature, we'll look at the range of different implementations that exist for storage virtualisation, as well as discussing the reasons why virtualisation exists at all.

Abstraction layers

Even at the lowest level, storage devices include a modicum of virtualisation (see the box *Disk controllers*, opposite top). In all areas of computing, we come across 'abstraction layers' – rigidly defined, standardised ways of allowing one device to communicate with another.

DAVID CARTWRIGHT explains how virtualisation makes storage systems easier for users by hiding some of the more complex underlying technology.

Konqueror provides application-level virtualisation for 'everyday' protocols such as FTP and HTTP.



Abstraction layers reduce the apparent complexity of computing devices. As technology becomes more complex, it's necessary to hide this complexity from the user – or even from the system manager for lower-level technologies. If we can keep the architecture of this stuff on a 'need to know' basis, then by definition we're hiding the stuff that the user doesn't need to know, which means they can concentrate on the bits they do need to understand.

Abstraction also allows us to 'forget' the underlying architecture of our kit, which allows us to build modular installations and evolve them as business requirements change. Imagine, for example, that you have a small fileservers with a single hard disk, connected over an IP network to a number of desktop PCs, and that each desktop PC has the server's disk mounted via NFS as directory `/acme/users`. We could replace the server hardware with a new model, and as long as it had the same IP address as the old one and advertised the disk to the network identically, the client PCs would require no reconfiguration. We could then move the single disk to an external array, and again the users wouldn't notice. Then we could replace the single disk with a four-disk RAID5 array, and again the users wouldn't notice the difference. (Not only this, but even the operating system probably wouldn't notice – the computer's BIOS would tell it it had a single disk, because that's what the RAID controller is telling the BIOS).

So we've talked around the benefits of virtualisation/abstraction – but what about the actual implementations? Let's look at the key areas in which you can implement virtualisation in a network, and go through the comparative benefits and drawbacks.

Array-level virtualisation

We've already mentioned RAID – Redundant Array of Inexpensive Disks – as a typical example of virtualisation. The common use of RAID is a RAID5 array – a setup in which you have three or more hard disks configured as a single virtual volume. The idea is simple: as the computer writes to

what it thinks is a single disk, the RAID controller actually writes information to multiple disks, in such a way that every item of data is represented one more than one disk. So should one disk fail, it can be replaced and its contents rebuilt using the information stored on the other disks in the array.

The main benefit of this approach, which uses a dedicated hardware adaptor card to handle the virtualisation, is that management of the RAID configuration is separate from management of the way disks are used by the computer. Instead of the computer operating system seeing every disk in every array, it only sees the 'virtual volumes' – which means the system manager is not overloaded with information in one place. The drawback is mainly the cost of the additional hardware required, but there are inexpensive entry-level options on the market (sub-£200 for SCSI RAID cards, sub-£100 for ATA).

Software RAID

The RAID concept discussed above assumes that the RAID handling is done in hardware. The same can be achieved without the need for a RAID card, and in fact Linux's implementation of 'software RAID' is widely regarded as one of the fastest.

Software RAID achieves the same abstraction as far as the operating system is concerned, as a clump of several disks looks like one. And although the RAID management is now part of the operating system setup, the RAID and normal disk management tasks are performed in separate configuration applications, so there's no extra confusion. The main loss is in performance, no matter how clever the RAID algorithm – instead of being performed on a separate adaptor, it's giving the computer's CPU and disk controllers more work to do.

Storage area networks

The next obvious step for storage virtualisation is to move the disks outside the computer. If you have a simple SCSI or Fibre Channel connection to an external disk array, there's really no difference between that and having internal disks, but if you go to the extent of putting some kind of multi-access control between the disks and the computer, you have a Storage Area Network, or SAN.

The most common type of SAN uses Fibre Channel connections, because they're (a) fast; (b) able to span much longer distances than SCSI chains; and (c) designed from

DISK CONTROLLERS

AT THE MOST BASIC LEVEL, THE HARD disks (IDE/ATA or SCSI) we put in our computers already have an element of virtualisation. These days, most disks have on-board controllers that handle the complex underlying electronic transmissions required to read from and write to the disk, and wrap them up in a developer-friendly command set. SCSI is an excellent example – if you're writing a program that talks directly to a SCSI device you literally have commands such as 'Rewind', 'Stop/Start Unit' and 'Prevent/Allow Media Removal'.

Given that one of Linux's strengths is its ability to run on any old PC we dig out of the cupboard and decide to make into a DNS or DHCP server, many of us will already be using another piece of abstraction. Many large hard disks have a switch you can enable to make them pretend to be smaller than they really are, so they'll work with older computers. All around the world are computers with, say, 40GB hard disks pretending to be 32GB ones, so that the ageing hardware can cope – and all this is achieved through one little switch on the disk's controller card.

“Abstraction layers reduce the apparent complexity of computing devices, allowing us to build modular installations.”

the ground-up to be part of a switched storage network. The computers and disks connect to a Fibre Channel switch, and it is this switch that decides how to route the information between the various endpoints.

Although the computer and the disk clearly need to be able to (a) know each other's address and (b) perform authentication to ensure that data security is preserved, this is handled by a low-level layer within the operating system of the computer, and as far as the main meat of the OS is concerned, it's just another disk that looks like it's connected directly to the computer.

The usefulness of a SAN is that the storage devices and the computers can be physically separate, and the SAN switching technology can provide elements of RAID-like functionality (eg by sending data that the computer is writing to the disk to two separate storage devices instead of, as the computer believes, just one). The downside of SAN operation is that by its nature it uses very high-speed networking technologies, which means that switching hardware is expensive. 

THE PURPOSE OF ABSTRACT INTERFACES

As Linux Format and Linux Pro are fond of repeating, standards are important...

IMAGINE WE HAVE TEN computers, all made by different companies and with their own proprietary storage technologies. Now imagine we have ten disk drives, made by ten different storage companies, and each with its own proprietary technology. When networked, for all the computers to communicate with

all the types of storage device would require 10 x 10 (=100) interfaces to be defined. Unless, that is, we decided to define a common standard by which computers would communicate with disk drives. All we'd need to do is write 20 interfaces: one between each computer and the standard

intercommunication mechanism, and one between each storage device and the standard intercommunication mechanism.

The benefits are magnified when we add a new make of computer or storage device. In the old, proprietary world, adding an 11th computer would require all ten storage guys to produce

interfaces between their kit and the new computer. In a world with a common interface standard, though, all that's needed is for the new computer manufacturer to make sure his device works with the standard intercommunication technology – as long as it does, it'll instantly work with all ten types of storage device.

STORAGE

« iSCSI (Internet Small Computer System Interface)

A new technology that's just starting to see the light of day is iSCSI. Again, this is a classic example of virtualisation, because you effectively have two endpoints that think they're connecting over a standard (in this case SCSI) protocol, whereas in fact there's an arbitrarily large IP network providing the link between the two points.

SCSI-connected devices work by sending SCSI commands back and forth over a multi-wire cable across distances of a few metres. iSCSI's benefit is that this inter-device distance can be extended almost without limit by capturing the commands as they leave the SCSI port of each device, bundling them up ('encapsulating' them) inside IP packets, and sending them over an IP network to the destined recipient. At the receiving end, the SCSI commands are unbundled from the IP packet and sent into the receiving device via its SCSI port. Such packet encapsulation is not a new concept at higher layers in the network (we've been encapsulating AppleTalk and IPX inside IP packets, and IP within X.25 packets, for years). What's unusual with iSCSI is that a low-level protocol (SCSI) is being encapsulated within a high-level protocol.

Fairly obviously, iSCSI counts as a storage virtualisation technology because the endpoints of an iSCSI connection have no idea that they're actually connected not via a piece of SCSI cable, but by an IP network. iSCSI's benefit is that it allows storage networks to extend out of the LAN and into the MAN or the WAN – so one could have a pair of storage arrays in different physical locations, both storing the same data and thus providing extremely high immunity to flood, fire or electrical problems at a single site. The downside of iSCSI is that for it to perform at all well, the connection

LINUX VIRTUAL FILESYSTEM

WE'VE ALREADY SAID THAT using our abstraction concept, we can unify all storage devices and computer systems via a single interface; the idea of the Linux Virtual Filesystem (VFS) is that instead of having a collection of file system formats (DOS, Ext2, etc) an interface is made between each of these systems and a general definition of a filesystem – VFS. So just as abstraction enables us to have less product-specific interfaces between devices, so the introduction of VFS enables us to have support for one general filesystem in the filesystem layer of the OS, instead of a dozen or more.

between the endpoints must be fast enough not for it to become a bottleneck to system performance.

OS-level file sharing

SERVER MASQUERADING

We've looked at virtualisation techniques that work either outside the computer or at a very low level. The other main area for virtualisation is within the operating system itself, providing the ability to share storage areas among heterogeneous systems.

Linux is particularly flexible in its ability to make disks and directories available in a form that other machines can access natively. NFS (the Network File System) is native to most Unix-like implementations, Linux included, but there's also Windows file sharing (via *Samba*), AppleShare (via kernel AppleTalk drivers and the Netatalk application), and NetWare-style sharing (using Martin Stovers' NetWare Emulation or the LinWare Daemon).

When a client computer attempts to connect to a server via a file sharing protocol that the client implicitly understands, as far as the client is concerned, it has received the information that confirms it's talking to a server that natively supports its file sharing protocols. After all, the server is sending all the same signals to the client that a native server would be using – it just happens that someone had to add some bits to the server in order to enable this to happen.

Client-side protocol implementation

Another way to share files in a heterogeneous network is to leave the server untouched and enable the client to work with the server's built-in network protocols. For example, as

DISTRIBUTED FILESYSTEM PROTOCOLS

As always, Linux is about choice...


BACK IN THE DAYS WHEN LINUX fitted on a few floppies, it wasn't unusual to download just the minimal installer-boot floppy set (usually three or four disks) and then mount the archive over the Internet via NFS. The installer would then quite happily install the rest of the software via this NFS-mounted disk, though given that it would have been the early 1990s, the speed of the average Internet connection would make it an interminably slow process.

NFS was, however, intended for use within a corporate network – specifically a LAN. Although the protocols were modified over the years to make them more efficient (and therefore faster over wide area connections). For truly

distributed file system operation, there are some alternatives that were built with this type of operation much more in mind than was the case with NFS.

Coda (www.coda.cs.cmu.edu/) is a distributed filesystem project that has been around since the late 1980s, and which was designed not just as a way of sharing files but with flexibility in mind. So it includes the ability to work offline with local copies of files (essential for laptop users), it can vary its behaviour depending on the available bandwidth, and it can replicate data between servers and cache documents in order to keep load times low.

AFS is another popular distributed



Microsoft's Services for Unix allows Windows systems to work with Unix-style protocols such as NFS.

file system among Unix-style systems. It started as the Andrew File System project at Carnegie Mellon University, but is now managed by IBM (who took over Transarc, the original commercial maintainer). Like Coda, the fact that it was designed from day one as a distributed file system means

that it will replicate data in multiple locations for fast access, and IBM being IBM, they maintain OpenAFS as an Open Source project that you can access at <http://oss.software.ibm.com/developerworks/opensource/>. The main page for OpenAFS is available at www.openafs.org/

well as being a Windows-style fileserver, *Samba* can also act as a Windows client – so a Linux machine can mount a Windows fileserver so long as it has been given the right driver software. Similarly, a Mac or a PC can mount a Unix-native NFS volume, given the right optional software.

Client-side protocol implementation is far less ‘virtual’ than its server-side sibling simply because the client needs to understand something about the server’s capabilities. This said, the user doesn’t necessarily need to know that their computer has been given some magical new abilities – as far as they’re concerned, it’s just a directory on their computer that seems to work a bit slower than the others for some reason.

Application-level virtualisation

OS-level virtualisation works at a higher level than the disk-level alternatives such as iSCSI and RAID. There’s a final level of abstraction that sits on top of all of the OS-level shares and adds still more protocols to the mix – application-level abstraction.

Application-level storage virtualisation is typified by *Konqueror*, the all-in-one file browser that many Linux users (KDE users in particular) will no doubt be familiar with. In everyday use, it’s a simple filesystem navigator: what makes it unusual is that when you fire it up, the current location being browsed is described not as a directory name but as a Uniform Resource Identifier, or URI. So it might default to something like `file://users/david`, which corresponds to user David’s home directory on the local machine. Having read everything we’ve already said, you’re no doubt already thinking that this supposedly local, single directory could well be mounted via NFS from another server, which in turn mounts it via Windows filesharing from a Windows 2000 box, which is in fact connected via iSCSI to a storage vault, which contains a RAID5 array formed of seven disks. And this could well be the case – that’s the power of virtualisation.

We can, however, change the protocol we use to get to our file. We could, for instance, change the location to `ftp://src.doc.ic.ac.uk/rfc/rfc-index.txt`, which would enable use to browse the FTP archive at Imperial College’s SUNSite installation – in this case, though, *Konqueror* is having to do some work to initiate FTP connections and issue FTP commands behind the scenes. This is because the operating system isn’t able to hide this stuff from us – the virtualisation is being done at the application level. The same applies if we change the protocol to `http://...` – the application (or, more likely, the system libraries it sits upon) is having to do work because the operating system doesn’t abstract that far up.

Of course, this type of abstraction fills the final gaps in the virtualisation pictures: it allows the user to work on the principle that the file they want is just a file that exists somewhere, without having to worry about how to actually go about fetching it.

Summary

As we said at the beginning of this feature, storage virtualisation means different things to different people. In reality, though, in many cases we find ourselves using it

UNIFICATION ISN’T NECESSARILY GOOD

Standards may be important, but it’s key to know when they are appropriate.

WE’VE SAID ELSEWHERE THAT THE KEY benefit of virtualisation (or abstraction) is that it allows us to interface to things via generic interfaces – a ‘write once, use many times’ concept for any programs we write that use storage devices.

The downside of this approach is that in most cases, the more general you make something, the less efficient it becomes. This is because in order to bring a number of systems under one interface, it’s usually necessary to eliminate some of the esoteric little features that enable one or two specific devices to work optimally; after all, if you have one device with a special performance-enhancing widget, you can’t represent that in a generic interface that’s supposed to work with all types of device.

In fact, the SCSI people have already acknowledged this issue. Although SCSI, and particularly SCSI-2, is a very generic, ‘works-with-anything’ definition, with the new SCSI-3 standards, we actually have not one set of commands but one set for each type of device (tape drives, disk drives, etc). The reason is simple: the world is craving faster and faster storage devices and the most sensible way to enable the various types of device (which, let’s face it, are radically physically different from each other in most cases – for example, tape and disk drives) was to provide specific command sets for each. There’s still a core command set, but in the SCSI-3 standard you’ll also find a set of specific stuff for each device type.

“As networks get faster, both in the LAN and the wider area, concepts such as iSCSI will provide more and more flexibility.”

without even knowing; not only this, but we often use several different incarnations of virtualisation at the same time. It’s not unusual to have a RAID filesystem mounted via NFS to another machine, for instance, or to be browsing files in assorted locations via direct disk access, FTP and HTTP all in the same *Konqueror* window.

As networks get faster, both in the LAN and the wide area, newer concepts such as iSCSI will grow in popularity and provide more and more flexibility for system managers with esoteric storage requirements. We’re already used to networked filesystems where the system manager decides where everything lives and controls access to it. The next leap is to truly distributed filesystems where the infrastructure decides the most efficient location for each item (and can decide to keep cache copies in multiple places should it deem this necessary) – and in reality, many of us have already started to take that leap. ■

Mounting a Windows fileshare under Red Hat Linux 9.0.

```

root@blacksun:~# mount -t smbfs -o iocharset=utf8 //192.168.1.27/81
aread /a/mnt/smbfs-81
2076: session request to 192.168.1.27 failed (file/directory not present)
2077: session request to 192.168.1.27 failed (file/directory not present)
Password:
root@blacksun:~# df -h
Filesystem            1k-blocks    Used Available Use% Mounted on
/dev/hda1              617768      161348    456420   26% /
tmpfs                  131064      0         131064   0% /dev/shm
/dev/ram              10019072    10019072  0         100% /boot
root@blacksun:~# df -h
Filesystem            1k-blocks    Used Available Use% Mounted on
/dev/ram              10019072    10019072  0         100% /boot
/dev/hda1              617768      161348    456420   26% /
/dev/sda1              10019072    10019072  0         100% /mnt/smbfs-81

```

OPEN MINDS



CLUSTER or FAIL

Why high availability doesn't have to mean high anxiety

System downtime – either on desktops or in the server room – means lost productivity, lost opportunities, and ultimately lost revenue. Despite a glimmer of light hovering at the end of the current IT tunnel, not a single CTO in the world is willing to take risks in the current environment, which means systems need to be more locked down than ever.

With even 99 per cent uptime working out as three-and-a-half days of unplanned downtime each year, there's a lot of pressure on computing personnel to keep things running no matter what. Of course, the key problem with high-availability systems is that when they work, no one notices that the technology is there. Furthermore, the only way to properly test how resistant to downtime a network is to have a *real* disaster hit – not a scenario that's easily tested!

Given these inherent obstacles, how can companies suitably protect themselves from the threat of unplanned downtime? More importantly, how can this be achieved without the outrageous expense often associated with high-availability solutions?

Tried and trusted

One of the most commonly utilised ways to get a list of solutions for your own company is to examine what others are doing, and it's no different with high-availability. Looking

Keeping your computers up 100 per cent of the time should be a business reality in this day and age, but what's the most economical way to achieve it? PAUL HUDSON finds out...

around, you'll find that the BBC, BT Syntegra, Allied Irish Bank, Aston Martin, and others all use the same system for their high-availability solution, which should be a powerful tip-off that someone, somewhere, is doing something right.

Linux Pro tracked down that "someone", and found Open Minds High Availability Solutions, a company originally established 13 years ago as a Unix consulting firm. Since it was founded, Open Minds has become the leading UK solutions provider for *SteelEye LifeKeeper*, one of the most successful high-availability (HA) products that is available on the market right now, and the company has been deploying solutions based around this technology for some years now with a great deal of success.

The key to *LifeKeeper's* effectiveness lies in its flexibility – it can take a 32-way HA cluster in its stride with no requirement for application pairing, and supports a wide variety of Linux distros, including Red Hat 7 and 8, Red Hat Enterprise Linux AS 2.1, and United Linux-based distros. This is enhanced further using *LifeKeeper Application Recovery Kits (ARKs)*, which are currently available for DB2, Oracle, Apache, Samba, MySQL, and other popular applications. There is also a *Generic Recovery Kit* that can be easily customised to support nearly any application.

LifeKeeper can replicate data over a WAN or LAN, providing off-site replication and even more reliable disaster

recovery. Of course, although *LifeKeeper* is user-friendly to maintain, it takes experienced veterans in the field of clustering and disaster proofing to be able to make the most of the software, and this is where Open Minds comes in.

With hundreds of satisfied customers, the Open Minds ethos of 'service availability rather than hardware availability' is clearly hitting home. This is backed up by Open Minds' experience with both Red Hat and SUSE Linux, which means it is able to put together the optimal solution for each problem, and also means that it is able to provide clustering for each component of the whole system if necessary, as opposed to just some parts. For software developers, Open Minds is also able to work directly with developers to help increase the resilience of their applications without needing to modify the application or make it cluster-aware. This is usually a much cheaper way of making an application highly available than using a cluster API.

Why you need clustering

As HA software and hardware becomes more commoditised on the marketplace, the prices have tumbled a surprisingly long way – there's really not much excuse now for unplanned downtime simply because it costs so little to get started. The most simple clustering type is known as 'heartbeat', because one machine monitors another machine to make sure it's working. This is often considered to be the best on the grounds that it is so simple – there's a rule of thumb in HA circles that "complexity is the enemy of availability", and this is where the heartbeat principle wins out.

Some products, particularly from big-name vendors such as Oracle, come with their own high-availability solutions, and they are certainly powerful and well-suited to their individual tasks. However, you start to run into problems when you have three or four of these bespoke high-availability packages running – it's a waste of resources, it takes more staff training to administer, and it's almost certainly going to work out more expensive than if you were to use an all-in-one package such as *LifeKeeper*.

As one of the most popular clustering solutions available (with users such as Sony, Nokia, Delta Airlines, and Ernst & Young), *LifeKeeper* certainly has a convincing track record. However, how can you use it in your business?

Types of backup

Put as simply as possible, there are four types of backup system: standard backup, cold standby, cold standby with replication, and hot standby. Standard backup is what many companies still rely on today – you have a server with a removable backup device such as a tape drive, and when disaster strikes, you simply fix the computer and restore all the data. Unsurprisingly, this is the slowest way to handle backup and restore, and also has the largest potential for loss of data.

Cold standby is where you have a spare machine of the same configuration as your original machine ready to take over in the event of a problem; however, it doesn't have a live copy of the data (hence 'cold' standby). This is faster than a standard backup solution, but you still need to go through the hassle of performing backup and restore.

Hot standby is where a live copy of the data is stored on a replica server, waiting for disaster to strike. When things do

Here we can see *LifeKeeper* managing a cluster consisting of two nodes (d1 and d2), which are configured in an active/active configuration for numerous services.



LIFEKEEPER IS A TRADEMARK OF STEELEVE TECHNOLOGY INC

go wrong, the application is immediately recovered on the backup server using the latest copy of the data, effectively losing as little data and time as possible. Hot standby systems are designed to take full load within minutes (if not seconds) of the primary server going down, automatically rerouting clients from the downed server across.

The cold standby with replication option modifies the cold standby paradigm so that the backup server does have a live copy of the data, and therefore works more like a hot standby server. However, the replica server is only able to recover from hardware failure, and doesn't have any way to recover from application failure.

Each one of these backup types comes with its own advantages and unique requirements during deployment, which means that you need to be careful when making

“99 per cent uptime works out as 3¹/₂ days unplanned downtime each year, so there's pressure to keep things running...”

your choice. There are, unsurprisingly, a number of hidden pitfalls in there, which is why most companies prefer to have experts like Open Minds design and deploy their solutions for them, rather than attempt it in-house.

Peace of mind

Clustering technology may appear complex at first glance, but that's no reason to ignore it, particularly because end-to-end solutions can be provided by Open Minds. While the peace of mind factor cannot be underestimated, high-availability solutions are now so commoditised that you are simply at a competitive disadvantage if you do not deploy a trusted disaster-recovery solution in case the worst happens.

To make the choice even easier, you can even arrange for an Open Minds team to come to your office to demonstrate what they can do for your systems, as well as leaving you with a 30-day trial of the *LifeKeeper* software for you to try out at your own pace. This is really as easy as it gets, and won't cost you a penny either – is there any other excuse holding you back? 



◀◀ BEING OPEN-MINDED

Linux Pro spoke to **SHOBANA PATEL** at Open Minds about the need for a reliable clustering solution and what Open Minds brings to the market

LINUX PRO: How long has Open Minds been developing its clustering and high-availability solutions?

SHOBANA PATEL: Open Minds has been working on high-availability and disaster recovery for three years now. The high availability part of the company started as a result of a group of individuals who had been working with NCR, protecting on high-end systems. We project managed large corporate accounts from banks to haulage companies, to help them reduce single points of failure in their IT infrastructure. At that time, we were fortunate enough to be able to recognise that protection of systems was not just something large corporates needed, but mid-range and even some small users would need.

With this in mind, our aim with Open Minds High Availability Solutions was to make industrial strength high-availability affordable to a tier of the market that would not previously have considered it.

We incorporated a number of elements to our solutions that have removed the traditional barriers to entry of clustering and high availability. We have reduced the cost to the user, removed the requirement for shared storage and eliminated the need to modify applications to make them cluster-aware. By removing these requirements we have made high availability clustering and disaster recovery accessible to a whole new tier of the market.

Linux is a way of bringing the cost savings of Intel-based hardware to the solution. The increased efficiency and reliability of the hardware make this an attractive option. We have predominantly used the HP ProLiant hardware because of its cost and resilience.

We use a combination of Open Source and licensed software. Using Open Source where it does not compromise the solution, using licensed software where we feel that it provides the better option for the user. In our business, reliability, maintainability and resilience are the main considerations. We work to maximise service availability to the end-user.

LXP: And now the question that we always ask: how important is Linux to your plans?

SP: We were one of the first companies to focus specifically on Linux in the enterprise. We have worked hard with vendors like HP and IBM to get the message that Linux is enterprise-ready. That message wasn't being heard very clearly three years ago. Initially, when we went into customers we had to 'sell' the idea of Open Source and the fact that there were alternatives to Windows. Linux was important to us then, and after three years of hard work getting the message across and having proven time and time again that an industrial-strength solution can be created using Linux, we are not going to stop now.

The Linux argument is simply too strong for it to just go away. There are important drivers in the market that mean that corporations over the world are considering moving parts of their infrastructure onto Linux, and high availability



“We have reduced the cost of HA to the user, removed the requirement for shared storage and also eliminated the need to modify apps to make them cluster-aware”
SHOBANA PATEL, MD OPEN MINDS

is often an essential part of these plans. Our experience is that the high availability parts of the solution cement the components of the whole infrastructure together and strengthen it, but they need specialist experience to put it all together. The Linux market is growing up and we are well positioned to be a part of this growth.

LXP: Many companies tell us how pleased they are about Linux's lower licensing costs. How much of a factor is cost when you consider using Linux?

SP: Cost is a factor at the initial stages, but the functionality has to be there. Unless the products are suitable for their intended purpose, there is no solution – no matter how cheap they are.


Then, when a solution is seen as fit for purpose, the consequences of buying it and integrating it within the infrastructure is assessed. The solution has to be proven to work and maintainable. I have found that the Linux solutions we put in are scrutinised far more than their equivalent Windows counterparts.

GET SORTED


Four ways Open Minds can help keep your data safe

APACHE FAILOVER IN AN ACTIVE/ACTIVE ENVIRONMENT


When the Apache web server on server 1 fails, it will restart on server 2




In this scenario, shared storage is used. The Apache configuration and data files on SERVER1 are available to SERVER2, and vice versa for Samba. This keeps hardware costs to a minimum whilst giving big benefits in reliability.

APACHE FAILOVER IN A N+1 ENVIRONMENT

When there are multiple servers acting as a server farm, the n+1 solution is usually best. Here, one machine acts as the failover backup for all the machines. This is quite inexpensive, and there is no performance hit if a single failover occurs. This solution requires shared storage.

APACHE – REPLICATION OF DATA TO UPDATE WEBSITE (NON-SHARED STORAGE)

This picture demonstrates a hot standby with data replication. This system is usually a good bet for mid-range cost whilst still providing a good level of disaster-proofing. The servers shown here are separated by a LAN.

APACHE DISASTER RECOVERY OF A WEBSITE

This picture demonstrates a hot standby system with off-site data replication over a WAN.

As Linux establishes a foothold in an organisation, and is used for systems that are more critical and enterprise-ready, I believe that functionality rather than cost is the main deciding factor.

LXP: The MS-Blaster worm and SoBig virus woke many companies up to the fact that their current security implementations and policies are woefully inadequate. Are many companies still without sufficient disaster-recovery solutions?

SP: If you ask people if they have sufficient disaster recovery, then I think you'll find that most people will say they have. Our actual experience is that many companies

do not have adequate disaster recovery solutions, often because they do not realise the cost of downtime – for a simple example, imagine a small travel agents with five employees. For the last few years they have been using a bespoke application for booking and arranging holidays, which runs on a PC server in the broom cupboard. If a hard disk fails in the server, how long does it take to fix? Even a few hours downtime (or lost data) could cost thousands in terms of lost revenue and staff overtime. Then you need to factor in external consultancy just fix the problem and get things running again where they left off.

A hot standby solution can in comparison seem cheap, allowing maintenance to be scheduled and



OPEN MINDS

the failure to go unnoticed. The solution does not require expensive hardware (for instance, SCSI shared storage, SAN, NAS etc), as data replication is a viable and popular solution. A hot standby system can be implemented on any PC-based hardware that our notional travel agents might use, ranging from a low-end desktop to a high-end server with built-in hardware reliability.

LXP: Are companies becoming more educated about the need for system recovery?

SP: Companies are increasingly recognising that although backups have their place in providing a measure of damage limitation, their organisations' uptime needs cannot be met through a backup-and-restore strategy. Therefore, more and more companies are looking at next-generation recovery methods, these may be in the form of data replication solutions, or they may be in the form of maintaining a hot backup server.

This has been a learning process for most organisations brought about by the increasing dependence on IT

HIGH RECOMMENDATION?

To find out whether Open Minds' actions live up to its claims, *Linux Pro* spoke to GRAHAM SMITH, Head of IT Services, FTSE Group.

LINUX PRO: What does the FTSE Group do?

GRAHAM SMITH: FTSE Group is an independent company whose sole business is the creation and management of indices and associated data services. FTSE has no capital markets involvement. The company originated as a joint venture between the *Financial Times* and the London Stock Exchange.

FTSE serves thousands of clients in 77 countries worldwide, working directly with pension plan sponsors, investment banks, brokers, consultants, fund managers, stock exchanges and data providers. Two-and-a-half trillion US\$ of assets are under management using FTSE indices.

FTSE calculates over 65,000 indices daily, including more than 600 real-time indices. The best-known indices are the FTSE UK Series, which includes the FTSE 100. The fastest-growing index series is the FTSE All-World Index Series, a benchmark covering 48 countries.

LXP: That's a considerable number of indices to manage! What systems are you running at present to manage that magnitude of data?

GS: We have approximately 10 HP/UX servers, 10 Sun Solaris servers and around 10 Linux Servers. The systems are used to calculate the various indices and the data is used worldwide. We use various databases for index calculation of real-time, End of day and custom index products.

LXP: How important is your data to you?

GS: As we are working with the financial markets where information is critical, we cannot be complacent about the data.

Our systems work around the clock because we are driven by customer expectations. Our website provides market indices data for point in time and historical indices, averaging 120,000 page impressions per day.

LXP: What drew you to Open Minds?

GS: Open Minds were recommended to us by a colleague when we were looking at a Linux solution. We then spoke to Open Minds who we found very helpful throughout the evaluation

process. The quality of the solution and the strength of the experience drew us to Open Minds.

LXP: What products/solutions/training did you purchase?

GS: We chose the HP ProLiant Cluster running Linux Red Hat AS 2.1 with *SteelEye LifeKeeper*. The Linux cluster acts as a fileserver using *Samba* and *LifeKeeper* protects the *Samba* machine.

LifeKeeper provides the basic IP and file system failover and we purchased its *Samba Recovery Kit*. The *Samba Recovery Kit* provides a mechanism to recover protected *Samba* file and print shares from a failed primary server onto a backup server. *LifeKeeper* can detect failures either at the server level (via heartbeat) or resource level (by monitoring the *Samba* daemons, or IP resources) so that control of the *Samba* resources is transferred to a backup server. Other *Samba* services may coexist on a *LifeKeeper* server.

We could have used the second server as a web server to produce an active/active environment, but after careful consideration and consultation we decided that the system was too important to us and chose to opt for the active/standby configuration instead.

In addition to purchasing the software and hardware, we purchased implementation and training from Open Minds High Availability Solutions. We wanted to do as much of the implementation ourselves so we used it all as a learning experience. One of our network administrators put the hardware together when it arrived, and installed Linux and *Samba*. We then called in Open Minds to modify the *Samba* configuration to work with *LifeKeeper*, they implemented the *Samba* recovery and ran a *LifeKeeper* administrator training course to allow us to be able to administer the implementation.

We were impressed that the whole implementation and training took only two days. I would recommend combing the two because the way that Open Minds did this was to do the training first, then because we had the training we could understand what was happening while implementing the solution.

LXP: What are the *LifeKeeper* servers used for?

GS: The servers are used for all cross-department sharing of information. The information store is shared globally; it contains all the internal documents, external products historical store, and the products themselves. We expect storage requirements to grow to over 7TB.

LXP: How important are the systems?

GS: Financial markets measure down time in milliseconds – failure is not an option.

LXP: Why did you choose Linux clusters rather than HP or Sun?

GS: We actually looked at several options including windows file sharing, the Linux solution running *LifeKeeper for Samba* looked to be the most resilient and scalable. Reliability and availability are key concerns in the environment, but have so far been impressed with the reliability of the new equipment and software distributions.

LXP: What would you say is the key benefit to using Open Minds and its software?

GS: Without doubt, Open Minds' expertise in *LifeKeeper* and the HP ProLiant hardware.

LXP: What difference has using Open Minds made to your business?

GS: It has given FTSE the flexibility to look at different solutions without worrying about downtime. We can also perform upgrades and systems maintenance without being concerned about downtime. We also know that although our systems and components are working 24/7 that any failure will be resolved without intervention.

LXP: What would you say to companies and organisations that don't currently have a high-availability solution in place?

GS: From my experience, everything in a production environment has to be 24/7 plus. Companies like Open Minds help to take the pain away from the process of implementing the solution, whilst keeping costs to a minimum.

infrastructure, and made possible by the reduction in price of clustering software and hardware.

Disaster recovery is becoming more of a board level responsibility. As such, it is being taken seriously throughout the organisation, and systems recovery is an integral part of a disaster recovery strategy.

LXP: How can Open Minds help companies keep their data secure and system running?

SP: Data availability is often mistaken for service availability. Data availability is meaningless to the end-user unless they can access their applications. The end-user requires the service to be available. To provide service availability, Open Minds looks at every component required to deliver that service to the user, incorporating hardware components, operating system components and applications. To make these applications available to users we work to eliminate single points of failure.

Some of the tools we use to achieve this are data replication, offsite disaster recovery, remote failover,

“If a server hard disk fails, even a few hours downtime or lost data could cost thousands in terms of lost revenue and staff overtime.”



synchronous data replication and local failover. We have developed methods to do this for any Linux or Windows application including *Exchange, MS SQL 2000, Apache, Oracle, DB2, MySQL, PostgreSQL, IIS, File sharing, NFS, NAS, LVM, etc.*

LXP: Where do you see the high-availability market going in the next few years?

SP: Short-term, I can see an increasing demand for offsite recovery using real-time replication. Looking further ahead, I think that the trend towards server virtualisation will continue and high availability products will play an increasing part in bringing these concepts into fruition. ■

COPING WITH FAILURE

These three figures demonstrate how a disaster-proofed system handles itself even under the worst scenarios.






Fig1 shows the optimal system, with end-users connecting to the server array to get its information from server A.

In **Fig2**, server A has failed, and server B has kicked into action. End-users are still able to access all the resources as per normal, and server A is presumably being fixed at this point.

Fig3 demonstrates the worst case, where all hardware in one location has gone down. Here server C is required to start serving end-users directly, which it is able to do because it has been receiving data over the WAN connection.

As you can see from these illustrations, if there's a hardware failure or even a site-wide connectivity failure, this company is going to be able to carry on computing as if nothing had changed. Revenue lost? A grand total of Zero pounds and zero pence. It is these kinds of figures that make high-availability so sought-after.



HOSTING

WEB HOST FOCUS

Hostway

Hostway is enviably synonymous with the words 'web hosting', and its reputation for consistent quality at affordable prices is unequalled elsewhere, having been around for such a long time. Over the years, Hostway has developed a range of services that closely mirror most common requirements of companies. With so many hosts choosing to specialise either in managed hosting or virtual hosting, Hostway offers both, giving customers the most choice; and still offers an email-only service – which is rare – so there's a very full range available.

At the lowest end, a virtually hosted server can be set up for just £7.95 a month, or £79.95 for a year, which gives you 100MB of disk space, 1GB of transfer, Perl and PHP, as well as the usual bundle of email boxes. Crucially, this service is also protected by a firewall by default, which is usually an added extra with other companies. At the top end of virtual hosting, £89.95 per month (£809.95 per year) gets you five times the disk space, fifteen times the bandwidth, SSL access, and a dedicated tech support telephone number.

For companies that need more power than virtual hosting can offer, Hostway has a range of managed and unmanaged server options, including a "build your own" solution where

PAUL HUDSON finds out about the importance of Linux to one of the oldest and most successful hosting companies around.

you can pick and choose every aspect of the services and hardware you want. These machines all come with Red Hat 7.2 as standard, and prices start at £149 a month, or £1490 for a year. The absolute top-end prebuilt dedicated server comes with three 36GB hard disks preconfigured for RAID striping, two IP addresses, a dual CPU core, and is fully managed to boot – all for £879 a month, or £8790 a year.

Beyond prices, Hostway has always gone a long way to make sure it configure its machines as users would ideally like them. All Hostway's Linux machines, for example, come fully preconfigured to support PHP, Python, XML, WAP/WML, CGI/Perl, GCC, C++, TCL and SSH. Furthermore, the configuration and setup of each of these is supported by its tech support staff, which means that even first-time webmasters can get to grips with advanced web technologies.

A good example of how much detail Hostway put into its systems is its PHP system – it comes pre-configured with the latest release, and also allows clients to reconfigure it to their own requirements, thereby letting them maximise the performance of their site without needing to upgrade to a more expensive solution. Hostway is evaluating an upgrade to PHP 5 once it has been released and tested thoroughly.

WAY TO HOST...

Linux Pro spoke to Rob Lovell, the President of Hostway, about the company and its plans for the future.

LINUX PRO: What are your unique advantages?

ROB LOVELL: Firstly, we are one of the only companies to offer telephone support seven days a week – customers can always talk to a real person.

We also incorporate a very broad range of features into our hosting plans, thereby reducing the 'hidden costs' customers have to face. We are upfront about our pricing structure, and continue to support new technologies. Customers also benefit from the added security of a truly global hosting company. Hostway has offices in Europe, North America, Asia and Australasia, which brings financial and technical stability and innovation.

Finally, our range of hosting plans and custom dedicated solutions allows Hostway to offer customers a real choice in terms of their platform,



Rob Lovell

“Linux is key to the operation of Hostway as a hosting provider, and underpins all of our in-house systems, including the SiteControl 3 Control Panel!”

ROB LOVELL, PRESIDENT – HOSTWAY

scripting, monetary and strategic decision-making.

LXP: What kind of hardware do you deploy,

and what can you tell our readers about Hostway's hardware plans for the future?

RL: We know that only the highest-rated

CUSTOMER SERVICE ON THE INFORMATION HIGHWAY

To find out how Hostway meets the needs of its customers, *Linux Pro* spoke to one of them – Aspin Interactive (www.aspininteractive.com)...

ASPIN INTERACTIVE DESIGNS, DEVELOPS and implements a wide range of web-enabled apps, such as website content management software, eCommerce systems, web integration services with business applications and many other web based products. We have worked with Hostway for two years and grown the relationship from shared Windows hosting to developing long-term hosting and security strategies for specialist projects.

One of the key factors in the relationship is our mutual development of Linux-based solutions. Hostway's experience in providing both shared and dedicated Linux products has enabled us to bring applications to market very quickly. We have recently launched two online services using Red Hat and eSmith.

During the decision-making process Hostway took a real interest. After detailed discussion of our requirements, Hostway came back to us with a range of tailored, targeted options which enabled us to put together the ideal solution for any requirement. It was Hostway's willingness to listen that appealed, and made us feel like more than just a number in its customer database.

The first major project we have worked together on is an extranet project for



Riverford Organic Vegetables, a home-delivery box scheme operated in the South of England. We have developed a complete web-based franchise management system that allows the franchisees to manage customers, orders, and accounts plus create vital business reports entirely through their web browser. The system also integrates with Sage Line 100 through another one of our Linux-based remote servers at Riverford's main farm in Devon. As far as we are aware, there is no other web-based business management systems like this online today in the UK. This is soon to be fully integrated with a new public website for Riverford which will provide additional online services to customers and enable further automation of the entire customer service and delivery process. The

website extranet is completely written in PHP with MySQL and runs on a dedicated server located in Hostway's Docklands datacentre.

The reason for choosing Linux for the development platform and Hostway for the hosting was simple: we needed reliability and performance at a good price. Linux gives us a stable environment, and combined with PHP and MySQL we have a low-cost development environment that has the performance needed for large-scale access.

Hostway's high level of customer support and monitoring provides us with a service that ensures excellent performance and continuous uptime. We also have peace of mind with regards to security, as Hostway has made a major investment in both firewall services and onsite physical security precautions.

Finally, we have made significant cost savings by investing in Open Source technologies with regards to future licensing. The project overheads are reduced as hardware systems are easy to set-up and manage, as well as helping to lower our development costs.

We are now working together with Hostway to develop the hosting project further with redundancy and disaster recovery solutions.

Unparalleled networking

Situated just over a mile from LINX, the main Internet gateway in the UK, Hostway has 255MB bandwidth capacity over two dedicated lines provided by separated backbones. To cope with the growing capacity, they have even more capacity waiting in the wings to take up the slack as necessary.

Most people don't realise just how big Hostway is – as a

worldwide company it has over 300,000 clients at the time of writing spread across its datacentres in the UK, the Netherlands, Korea, and the US, all of which are configured to survive system failure and still run at full capacity. In the UK, the Hostway offices are in the same building as its datacentre, which means that its support engineers have quick and easy access if there are any problems – not a feat many host companies can match! ■

components fully serve the needs of a flexible, reliable and demanding web host. Therefore, our hardware consists of servers built in-house, fitted with Intel Processors and motherboards, Ultra320 7th Generation IBM SCSI Hard Drives, and Adaptec SCSI and RAID controllers, to provide optimal operating conditions for our clients' sites.

The main advantage of building our hardware in-house is that we can provide significantly higher SLAs on response times and hardware fixes than third-party hardware suppliers. We guarantee a two-hour maximum response time, and comprehensive stocks of spares onsite ensures we meet these levels.

In the future, we plan to expand our datacentre and connectivity as our customer base grows.

LXP: What is your key target market?

RL: Our primary target market for Linux virtual and dedicated hosting consists of developers and designers who enjoy and prefer to use an Open Source environment for their Internet applications and websites.

LXP: How much does Linux figure into your hosting plans?

RL: Linux is key to the operation of Hostway as a hosting provider, and underpins all of our in-house systems, including the *SiteControl 3* Control Panel.

We utilise a custom version of the Red Hat Linux OS and Kernel, with *Apache*, *MySQL*, *Python* and *Zope*, for internal CRM and networking. Whilst we also offer Windows and ColdFusionMX hosting plans, we are dedicated in our goal to bring Linux to an ever-widening audience of designers and developers.

LXP: Where do you see the hosting market being in three years time?

RL: Studies have shown that the hosting market is becoming more consolidated, with only the truly global companies able to expand in the marketplace. Hostway has made recent acquisitions in Amsterdam, Vancouver and Sydney, allowing us to offer a greater range of products to an ever-widening market.

Also, as a higher proportion of residential homes come online with the rollout of broadband and lower connection prices, we are seeing a higher percentage of customers starting out in web design. Subsequently, we have introduced our own proprietary website tool, *SiteBuilder*, which provides a user-friendly, comprehensive system for easy site creation and maintenance.